

Universidad de Burgos

Big Data Reference Architecture
for Industry 4.0

Including Economic and Ethical Implications

JAN STROHSCHNEIN

THESIS DIRECTORS

PROF. DR. ANA MARÍA LARA PALMA

PROF. DR. JOAQUIN ANTONIO PACHECO BONROSTRO

PROF. DR. HEIDE FAESKORN-WOYKE

JANUARY 2021

Abstract

The development of technology is getting faster and faster. In 1965 Moore predicted that the number of transistors in a computer chip and therefore its processing power would double every year for the next ten years. However, newer technologies evolve even faster, and recently a new term was born to describe them: “exponential technologies”. Technologies such as big data, the internet of things, or artificial intelligence double in power or processing speed every year, while their cost halves, often also complementing or catalyzing each other. The rapid progress in Industry 4.0 is achieved through the combination of exponential technologies and further innovations in several fields, e.g., manufacturing, big data, and artificial intelligence.

The thesis at hand motivates the need for a Big Data reference architecture to apply Artificial Intelligence in Industry 4.0 and presents a cognitive architecture for artificial intelligence – CAAI – as a possible solution, which is especially suited for the challenges of small and medium-sized enterprises.

Initially, the introduction of Artificial Intelligence, Big Data, and Industry 4.0 establishes common ground and explains why the intersection of those exponential technologies holds enormous economic potential. Subsequently, the economic and ethical implications of AI, Big Data, and Industry 4.0 are examined. These technologies show the potential to revolutionize the global economy and our working lives. Therefore, many national research programs and business ventures emerge in a race of economic powers for global competitive advantages, even though the ethical consequences are not yet completely foreseeable. Workers fear the loss of their jobs and the independence of their work. However, employee support is one of the most important criteria for a successful implementation. Therefore an ethical introduction process is presented, which is including and assisting rather than overburdening the employees.

Large corporations develop most Industry 4.0 technologies and standards. The “Industry 4.0 Questionnaire for SMEs” was conducted to gain better insight into smaller and medium-sized companies’ requirements and needs. The majority of participants stated, amongst other things, that they call for assistance to formulate an Industry 4.0 strategy and evaluate the available Industry 4.0 technologies. They need standards or best practices, which unfortunately are not yet available.

Reference architectures are introduced as they represent best practices for designing a software system and support companies during implementation. An evaluation of the proposed reference architectures for Industry 4.0 shows the possibilities but also the existing shortcomings.

Subsequently, the CAAI architecture for the application of AI in I4.0 production systems is presented. The provided set of standard architecture building blocks eases the implementation effort and assists companies in the Industry 4.0 introduction. Different use cases demonstrate the applicability of the architecture and the following evaluation verifies the capabilities and functionality of the CAAI architecture.

The resulting modular architecture can be implemented on existing IT resources or through various cloud computing providers. Thus, a small or medium-sized enterprise can use the architecture to optimize production processes even if their budget does not allow them to buy the required hardware outright.

Wider adoption of the CAAI architecture will further increase its usefulness as it is an open-source project and users can share their created modules with the community which decreases future implementation effort for all users.

Acknowledgements

Throughout the writing of this thesis, I have received a great deal of support and assistance.

First of all, I would like to express my sincere gratitude to my supervisors Prof. Dr. Ana María Lara Palma and Prof. Dr. Joaquin Antonio Pacheco Bonrostro as well as the University of Burgos for making this doctorate thesis possible and welcoming me with open arms. The guidance and fruitful discussions were important for the completion of this thesis and the valuable suggestions improved my work.

I would also like to thank Prof. Dr. Heide Faeskorn-Woyke, who has been my mentor for many years. She always encouraged me to pursue the research topics I was passionate about and has been a great source of support.

The interdisciplinary work in the KOARCH project and the constant exchange of ideas and discussions were also extremely valuable for me and my work. Thus, I would like to thank everyone in the project and especially Andreas Bunte and Andreas Fischbach, my fellow doctoral students.

Further gratitude goes to my family for their support during my academic education, especially my father, who read all my work and provided valuable feedback with an outside perspective.

Finally, I would like to thank Maren, my love, for her patience, encouragement, and support throughout the thesis.

Publications

Significant parts of this work are based on previously published material or material in preparation for publication during this thesis's writing.

The thesis was written while working on the “KOARCH” project, which has the goal to support the introduction of Industry 4.0, especially in small and medium-sized enterprises. KOARCH is a joined project between TH OWL in Lemgo and TH Cologne and supported by the German Federal Ministry of Education and Research (BMBF).

The following documents are of significance:

- Detecting Emotions in Social Media - A technological challenge to enhance youngest behavior (Strohschein et al., 2019)
- Evaluation of Cognitive Architectures for Cyber-Physical Production Systems (Bunte, Fischbach, et al., 2019)
- Employee Technology Acceptance of Industry 4.0 in SMEs (Strohschein, Lara-Palma, et al., 2020)
- CAAI - A Cognitive Architecture to Introduce Artificial Intelligence in Cyber-Physical Production Systems (Fischbach et al., 2020)
- Cognitive Capabilities for the CAAI in Cyber-physical Production Systems (Strohschein, Fischbach, et al., 2020)

All these works have been extended, rewritten, and restructured in some form before their inclusion into this thesis. The rewriting and restructuring aim to increase the clarity, unify notation and terminology, and put the described contributions in this thesis's overall context. In most cases, this includes a revision or extension of the content. If not otherwise specified, these publications are, in large parts, a contribution of the author of this thesis.

This is in no way meant to reduce the contributions of the other authors. They inspired the presented ideas, helped improve the writing quality, and supported the architecture implementation.

For the sake of clarity, these remarks are repeated at the start of the respective sections or chapters.

TABLE OF CONTENTS

| | | |
|-----|--|----|
| 1 | Introduction..... | 11 |
| 1.1 | Big Data | 12 |
| 1.2 | Artificial Intelligence..... | 16 |
| 1.3 | Industry 4.0..... | 20 |
| 1.4 | Conclusion | 25 |
| 2 | Economic implications..... | 26 |
| 2.1 | Brand Value and Market capitalization | 26 |
| 2.2 | GDP and Productivity | 28 |
| 2.3 | Government Initiatives..... | 31 |
| 2.4 | Conclusion | 34 |
| 3 | Ethical implications..... | 35 |
| 3.1 | Influence of AI and I4.0 on the global labor market..... | 35 |
| 3.2 | The new Digital Divide..... | 37 |
| 3.3 | Quality of work and work alienation | 37 |
| 3.4 | Accountability and trust in new technologies..... | 41 |
| 3.5 | Designing the relationship between humans and machines | 43 |
| 3.6 | The ethical Industry 4.0 introduction process..... | 44 |
| 3.7 | Conclusion | 46 |
| 4 | Industry 4.0 Questionnaire for SMEs | 47 |
| 4.1 | Industry 4.0 Self-Assessment | 48 |
| 4.2 | Technology Acceptance Model | 48 |
| 4.3 | Methodology | 48 |
| 4.4 | Results | 50 |
| 4.5 | Discussion | 56 |
| 4.6 | Conclusion | 58 |
| 5 | Reference Architectures for Industry 4.0 | 59 |
| 5.1 | Reference Architecture | 59 |
| 5.2 | Use Cases and Requirements | 63 |
| 5.3 | Overview of reference architectures for Industry 4.0..... | 67 |
| 5.4 | Evaluation | 78 |
| 5.5 | Conclusion | 81 |
| 6 | CAAI - A Big Data Reference Architecture for I4.0 | 82 |
| 6.1 | Layers..... | 83 |

| | | |
|-----|--------------------------------|-----|
| 6.2 | Bus Structure | 84 |
| 6.3 | Big Data Platform | 86 |
| 6.4 | Standard Building blocks | 94 |
| 6.5 | Use Case Implementation | 111 |
| 6.6 | CAAI Evaluation | 132 |
| 6.7 | Conclusion | 134 |
| 7 | Future Research..... | 135 |
| 8 | Conclusion | 136 |

LIST OF FIGURES

| | |
|--|----|
| FIGURE 1 - MAPREDUCE EXAMPLE | 14 |
| FIGURE 2 - THE MAIN MODULES OF A KNOWLEDGE-BASED AGENT | 17 |
| FIGURE 3 - THE FOUR INDUSTRIAL REVOLUTIONS..... | 20 |
| FIGURE 4 - MOST VALUABLE BRANDS IN 2019 | 26 |
| FIGURE 5 - MARKET CAPITALIZATION WITH YEAR-TO-DATE DIFFERENCE | 27 |
| FIGURE 6 - SIMULATED ECONOMIC IMPACT OF AI, BASED ON (MANYIKA ET AL., 2018) | 28 |
| FIGURE 7 - ECONOMIC IMPACT OF AI INTRODUCTION PER REGION..... | 29 |
| FIGURE 8 - MODEL OF THE RELATIONSHIP BETWEEN MACHINE PACING AND TECHNOLOGICAL CHARACTERISTICS..... | 39 |
| FIGURE 9 - EMPLOYEE COMPETENCE MODEL..... | 45 |
| FIGURE 10 - TECHNOLOGY USAGE IN COMPANIES WITH UP TO 49 VS MORE THAN 49 EMPLOYEES. | 51 |
| FIGURE 11 - Q14 INDUSTRY 4.0 IMPLEMENTATION STATUS COMPARING COMPANIES WITH UP TO 49 EMPLOYEES AND COMPANIES WITH MORE THAN 50 EMPLOYEES | 52 |
| FIGURE 12 - INDUSTRY 4.0 INDICATORS COMPARING COMPANIES WITH UP TO 49 EMPLOYEES AND COMPANIES WITH MORE THAN 50 EMPLOYEES..... | 53 |
| FIGURE 13 - SYSTEMATIC TECHNOLOGY AND INNOVATION MANAGEMENT IN COMPANIES WITH UP TO 49 AND MORE THAN 50 EMPLOYEES. | 53 |
| FIGURE 14 - COMPARING COMPANIES THAT AGREE/DISAGREE TO Q8 (MEAN + STD.)..... | 54 |
| FIGURE 15 - H ₅ OVERVIEW WITH COMPANIES GROUPED BASED ON THEIR Q10 ANSWERS (MEAN + STD.) | 54 |
| FIGURE 16 - H ₆ OVERVIEW WITH COMPANIES GROUPED BASED ON THEIR Q10 ANSWERS (MEAN + STD.) | 55 |
| FIGURE 17 - Q15 I4.0 INDICATORS WITH COMPANIES GROUPED BY COUNTRY..... | 56 |
| FIGURE 18 - CPPS ARCHITECTURE | 60 |
| FIGURE 19 - VERSATILE PRODUCTION SYSTEM, PICTURE BY ANDREAS BUNTE | 64 |
| FIGURE 20 - CONCRETE SPREADING MACHINE WITH CASINGS | 65 |
| FIGURE 21 - RAMI4.0 OVERVIEW | 68 |
| FIGURE 22 - TYPES AND INSTANCES IN A LIFE CYCLE | 69 |
| FIGURE 23 - INDUSTRY 4.0 COMPONENT..... | 70 |
| FIGURE 24 - ADMINISTRATION SHELL COMPOSED OF RELEVANT SUB-MODELS..... | 71 |
| FIGURE 25 - INDUSTRIAL INTERNET REFERENCE ARCHITECTURE OVERVIEW..... | 72 |
| FIGURE 26 - 5C ARCHITECTURE | 74 |
| FIGURE 27 - SOAR ARCHITECTURE OVERVIEW..... | 75 |
| FIGURE 28 - ACT-R ARCHITECTURE OVERVIEW..... | 77 |
| FIGURE 29 - GAP IN PROPOSED REFERENCE ARCHITECTURES | 81 |
| FIGURE 30 - CAAI ARCHITECTURE OVERVIEW | 83 |

| | |
|---|-----|
| FIGURE 31 - COMPARISON: VIRTUAL MACHINE AND CONTAINER | 87 |
| FIGURE 32 - DOCKER ARCHITECTURE OVERVIEW | 87 |
| FIGURE 33 - KUBERNETES ARCHITECTURE OVERVIEW | 89 |
| FIGURE 34 - MONOLITHIC APPLICATION AND MICROSERVICES | 90 |
| FIGURE 35 - SCALABILITY OF MICROSERVICES | 90 |
| FIGURE 36 - MESSAGING VIA PUBLISH / SUBSCRIBE..... | 91 |
| FIGURE 37 - CHAINING PRODUCERS AND CONSUMERS TO CREATE A PROCESSING PIPELINE..... | 92 |
| FIGURE 38 - COGNITIVE BIG DATA PLATFORM | 93 |
| FIGURE 39 - BUILDING BLOCK STRUCTURE | 94 |
| FIGURE 40 -INSTANTIATING A CAAI BUILDING BLOCK | 96 |
| FIGURE 41 - KAFKA WORKFLOW | 97 |
| FIGURE 42 - KAFKAPC CLIENT MODULE | 98 |
| FIGURE 43 - POSTGRES INITIALIZATION PHASE..... | 99 |
| FIGURE 44 - POSTGRES OPERATION PHASE | 99 |
| FIGURE 45 - POSTGRES MODULE | 100 |
| FIGURE 46 - POSTGRES MODULE PERSISTS INCOMING MESSAGES IN POSTGRESDB..... | 100 |
| FIGURE 47 - PLOT OF THE NEW PROCESS PARAMETER VALUE AND THE DECIDING ALGORITHM | 103 |
| FIGURE 48 - PLOT OF THE CPU RESOURCE USAGE PER ALGORITHM | 104 |
| FIGURE 49 - FASTAPI HMI OVERVIEW | 106 |
| FIGURE 50 – FASTAPI ROUTE DETAIL | 107 |
| FIGURE 51 - COGNITION COMMUNICATION OVERVIEW | 110 |
| FIGURE 52 - COGNITION BUILDING BLOCK | 111 |
| FIGURE 53 - CONCRETE CAAI ARCHITECTURE FOR EMOTION DETECTION USE CASE | 114 |
| FIGURE 54 - DATA COLLECTION AND ANALYSIS PIPELINE..... | 114 |
| FIGURE 55 - CONCRETE CAAI ARCHITECTURE FOR POPCORN PRODUCTION USE CASE | 119 |
| FIGURE 56 - POPCORN PRODUCTION MACHINE LEARNING PIPELINE | 120 |
| FIGURE 57 - CPU USAGE PER PRODUCTION CYCLE | 121 |
| FIGURE 58 - PREDICTION ERROR PER PRODUCTION CYCLE | 122 |
| FIGURE 59 - OBJECTIVE FUNCTION VALUE PER PRODUCTION CYCLE | 123 |
| FIGURE 60 - KUBERNETES DEPLOYMENT | 124 |
| FIGURE 61 - KUBERNETES JOB..... | 125 |
| FIGURE 62 - COGNITION DYNAMICALLY CREATES NEW PIPELINES | 126 |
| FIGURE 63 - DECLARATIVE GOAL DEFINITION | 126 |
| FIGURE 64 - CPU CONSUMPTION OF DIFFERENT OPTIMIZATION ALGORITHMS ON THE GROUND-TRUTH..... | 130 |
| FIGURE 65 - AGGREGATED PERFORMANCE RANKS FOR DIFFERENT OPTIMIZERS ON THE SIMULATIONS..... | 131 |

LIST OF TABLES

| | |
|---|-----|
| TABLE 1 - POSSIBLE COST REDUCTIONS THROUGH I4.0 | 30 |
| TABLE 2 - INDUSTRY 4.0 QUESTIONNAIRE FOR SMES HYPOTHESIS | 50 |
| TABLE 3 - INDUSTRY 4.0 QUESTIONNAIRE FOR SMES RESULTS | 57 |
| TABLE 4 – CONSOLIDATED REQUIREMENTS FOR CPPS REFERENCE ARCHITECTURES IN I4.0..... | 67 |
| TABLE 5 - EVALUATION OF THE REQUIREMENTS FOR PROPOSED REFERENCE ARCHITECTURES | 78 |
| TABLE 6 - COMPARISON OF CAAI BUS SYSTEMS..... | 85 |
| TABLE 7 - BUILDING BLOCK STRUCTURE AND PURPOSE..... | 95 |
| TABLE 8 - TWITTER USER FEATURES AND DESCRIPTIONS..... | 115 |
| TABLE 9 - TWEET ANALYSIS FEATURES AND DESCRIPTIONS | 116 |
| TABLE 10 - TWEET ANALYSIS MULTI-CLASS CLASSIFICATION EXAMPLE | 116 |
| TABLE 11 - AVERAGE EMOTIONS OF ALL USERS PER YEAR | 116 |

| | |
|---|-----|
| TABLE 12 – EMOTIONS OF USERS WITH LESS THAN 25% OF FOLLOWERS | 117 |
| TABLE 13 – EMOTIONS OF USERS WITH MORE THAN 75% OF FOLLOWERS | 117 |
| TABLE 14 - CHOSEN OPTIMIZERS WITH PARAMETER RANGES AND DEFAULT VALUES | 129 |
| TABLE 15 - EVALUATION OF THE REQUIREMENTS FOR CAAI..... | 132 |

CODE LISTINGS

| | |
|--|-----|
| CODE LISTING 1 - AVRO SCHEMA EXAMPLE | 93 |
| CODE LISTING 2 - PLOT SCHEMA..... | 101 |
| CODE LISTING 3 - FASTAPI CONFIGURATION | 105 |
| CODE LISTING 4 - KNOWLEDGE REPRESENTATION OF RANDOM FOREST | 108 |
| CODE LISTING 6 – KUBERNETES JOB DEFINITION | 125 |
| CODE LISTING 5 – COGNITION ALGORITHM SELECTION..... | 127 |

ACRONYMS

| Abbreviation | Long Form | Page |
|---------------------|--|-------------|
| AI | Artificial Intelligence | 11 |
| ALPAC | Automatic Language Processing Advisory Committee | 18 |
| API | Application Programming Interface | 90 |
| AWS | Amazon Web Services | 15 |
| BDP | Big Data Platform | 12 |
| CAAI | Cognitive Architecture for AI in CPPS | 12 |
| CL | Conceptual Layer | 85 |
| CPPS | Cyber-physical Production System | 12 |
| CPS | Cyber-physical System | 21 |
| DPL | Data Processing Layer | 85 |
| ERP | Enterprise Resource Planning | 22 |
| GCP | Google Cloud Platform | 15 |
| GDP | Gross Domestic Product | 28 |
| GDPR | General Data Protection Regulation | 32 |
| HMI | Human Machine Interface | 65 |
| HR | Human Resources | 44 |
| I4.0 | Industry 4.0 | 11 |
| IIC | Industrial Internet Consortium | 33 |
| IIoT | Industrial Internet of Things | 33 |
| IIRA | Industrial Internet Reference Architecture | 73 |
| IoT | Internet of Things | 11 |
| IT | Information Technology | 21 |
| MES | Manufacturing Execution System | 24 |
| MHACL | Mental Health Action Checklist | 39 |
| OPC UA | Open Platform Communications Unified Architecture | 64 |
| PLC | Programmable Logic Controller | 22 |
| QR | Quick Response | 25 |
| RA | Reference Architecture | 59 |
| RAMI 4.0 | Reference Architecture Model Industry 4.0 | 68 |
| RDD | Resilient Distributed Dataset | 14 |
| RFID | Radio-Frequency Identification | 22 |
| SME | Small and Medium-sized Enterprise | 11 |
| TAM | Technology Acceptance Model | 48 |
| VDMA | The German Mechanical Engineering Industry Association | 22 |
| VPS | Versatile Production System | 25 |

*Any sufficiently advanced technology
is indistinguishable from magic.*

*Sir Arthur Charles Clarke, Profiles of the Future:
An Inquiry into the Limits of the Possible (1973 Edition)*

1 INTRODUCTION

The development of technology is getting faster and faster. In 1965 Moore predicted that the number of transistors in a computer chip and therefore its processing power would double every year for the next ten years (Moore, 1965). This prediction would later become known as Moore's law and still holds true today. However, newer technologies evolve even faster and recently a new term was born to describe them: "exponential technologies". Technologies such as big data, the internet of things (IoT) or artificial intelligence (AI) double in power or processing speed every year, while their cost halves, often also complementing or catalyzing each other (Haupt, 2016; Michaelis, 2019).

The rapid progress in Industry 4.0 (I4.0) is achieved through the combination of exponential technologies and further innovations in several fields, e.g., manufacturing, big data and artificial intelligence (Deloitte Switzerland, 2015). While none of those technologies is revolutionary in itself, it is their fusion and the availability of bulk data, that enables the development of services that have not been possible so far (Drath & Horch, 2014).

Companies around the globe struggle to keep up with the fast-changing environment and adjust to the technological transformations. Unfortunately, small and medium-sized enterprises (SMEs) do not have the budget or man-power to build highly customized AI and big data solutions from scratch. It is therefore absolutely important to support them with an easy-to-use solution, which creates good results with little manual effort.

However, further automation and optimization will also change the way humans work. While companies still heavily rely on the practical knowledge of their workers and depend on their creativity, certain repetitive tasks will be automated and workers need to learn new skills to fulfill the requirements of emerging jobs. Those new jobs will call for a high degree of cooperation between humans and robots. Therefore, how workers experience the interactions with machines is crucial to ensure lasting worker productivity and the long term success of a company.

Organizations need assistance to successfully introduce I4.0, implement AI in I4.0 software systems and maintain or even improve their work culture. Therefore, the following contributions are made by the work at hand:

- Derive ethical guidelines, how the new technologies should be introduced to reassure workers and create a positive work environment that fosters technology acceptance.
- Conduct a questionnaire for a technological and cultural self-assessment of I4.0 readiness with manufacturing SMEs.
- Present a big data reference architecture to implement software systems for various use-cases and create technical guidelines for different scenarios.

These contributions determine the structure of the thesis at hand. Thus, the remainder of this chapter introduces big data, AI and I4.0 to establish common ground and explains why the combination of those technologies holds enormous potential.

The second chapter investigates the economic implications of the new technologies on the global economy, but also individual companies. Thus, it presents the possibilities for additional profits through emerging products or services as well as expected cost reductions. It also highlights the national research initiatives as economic powers race for global competitive advantages.

Even though the economic potential is huge it is not possible to precisely forecast the ethical implications. The third chapter examines the influence of big data, AI and I4.0 on the workforce overall, but also the individual worker and the changes towards the work environment and work itself. The findings are summarized and presented as ethical introduction process for I4.0.

I4.0 represents a particular challenge for SMEs. Therefore, chapter four introduces a questionnaire specifically tailored towards SMEs and their employees. The results suggest that companies need assistance with the introduction process to create an I4.0 strategy and design a positive human-machine relationship. Furthermore, the results show that SMEs require cost efficient solutions to be able to implement I4.0 as resources are scarce even though the potential for optimization is high.

A reference architecture can be used as guideline to reduce the implementation effort for an I4.0 software system. Thus, chapter five introduces reference architectures as cost efficient solution for software implementation and explains why it is necessary for modern cyberphysical production systems (CPPS) to include cognitive components. A catalogue of requirements, derived from I4.0 use cases, and the subsequent evaluation shows the possibilities and shortcomings of the existing reference architectures for implementing an I4.0 software system.

Chapter 6 presents a novel cognitive architecture for AI – CAAI – as an alternative to those existing reference architectures. The architecture overview explains the architecture itself and the technologies that compose the big data platform (BDP). It also introduces several of the prepared modules that reduce the implementation effort in a use case. Subsequently, several use cases demonstrate the applicability of the architecture and the chapter concludes with the CAAI evaluation.

Possibilities for future research are detailed in chapter 7, ranging from further development of the architecture as open-source software project, to collaborations with industry partners to implement the architecture in additional use cases, and accompanying the introduction process in a manufacturing company to derive and refine best practices for SMEs.

Chapter 8 concludes the thesis and summarizes the research and findings. Thus, the conclusion addresses each of the previously mentioned contributions and places those in the context of the work at hand.

1.1 BIG DATA

Several of the most valuable companies in the world are data-companies. This highlights the shift from physical to digital value creation. Even if a company creates physical goods, there is potential to optimize production processes, marketing activities as well as distribution networks. Not using big data in a highly contested field is not an option if a company wants to stay competitive.

The term “big data” was coined by Doug Laney in 2001, where he also introduced the “3 V’s of big data” as the different concerned dimensions (Laney, 2001). Laney and Beyer defined “big data” again in 2012 as “high volume, high velocity, and/or high variety information assets that require new forms of processing to enable enhanced decision making, insight discovery and process optimization” (Laney & Beyer, 2012).

Gantz and Reinsel created their own definition in 2011: “Big Data technologies describe a new generation of technologies and architectures, designed to economically extract value from very large volumes of a wide variety of data, by enabling high velocity capture, discovery and/or analysis” (Gantz & Reinsel, 2011).

Both definitions describe the growing amounts of data and the changes in data processing that are necessary to generate economic value. They characterize the different dimensions of data through the “V’s of big data”, which were originally introduced by Laney in 2001 as follows:

- **Volume:** the type and detail of data being collected. Before the explosion in computing power, businesses and governments collected data but had a challenging time storing what was collected. Today, the volume of data collected from consumers and by agencies continues to grow, but because of computing capacity, storage is no longer an issue. This means that firms and agencies no longer have a data problem but instead have a computing puzzle.
- **Velocity:** the speed at which data are collected. Data are no longer lagged. Instead, data are being collected in real time at incredibly fast rates.
- **Variety:** the types of data being collected. Whereas basic demographic data, attitudes and opinions, and possibly geographic information might have been collected in the past, today nearly anything and everything a consumer does online is being captured.

Since then many researchers extended those original V’s to fit their circumstances. Adrian Tole presented 5 V’s in 2013 and added value and veracity to assess if the data represents the truth and holds enough value to justify the costs of collecting, storing and analyzing the data (Tole, 2013). An extensive list of the different V’s with their origin and meaning was created by Professor Chuck Cartledge in 2016 (Cartledge, 2016).

However, bringing Big Data from theory to praxis required a shift in programming paradigms. In the early 2000s Google engineers reached data processing limits while building the Google search index. Scaling server hardware vertically, i.e., utilizing few but very powerful servers, was too costly. Thus, the engineers used commodity hardware, which was cheaper and easier to maintain. However, the server hardware was error prone when working on such massive data sets and the software was not robust to failures. Dean and Ghemawat developed the MapReduce programming model to overcome these obstacles (Dean & Ghemawat, 2004). The programming model splits huge workloads into manageable pieces, processes the smaller data chunks and finally aggregates the intermediate outcomes into a final result. Worker nodes can be easily added to the pool of available resources if more processing power is required. If parts of the processing system fail, it is easy to recover or restart only the affected chunks. A more complex software implementation is the tradeoff in favor of a system that is horizontally scalable and robust against failures.

Google’s MapReduce and Google File System papers inspired the creation of Hadoop¹ in 2006, the first software framework to implement distributed storage and data processing (Ghemawat et al., 2003). Figure 1 presents a MapReduce example based on the most recent Hadoop implementation (Apache Software Foundation, 2020).

¹ Apache Hadoop Website: <https://hadoop.apache.org/> retrieved 08-11-2020

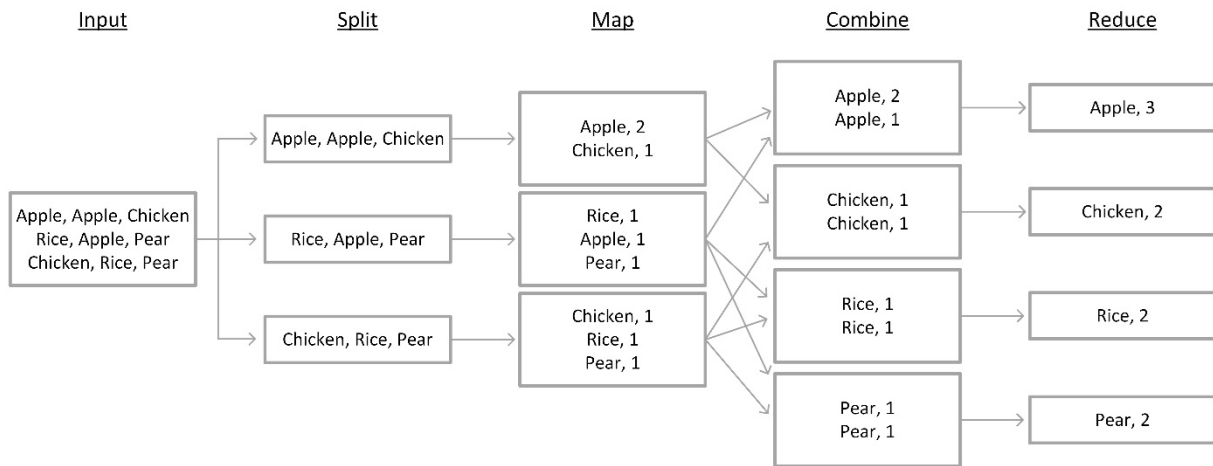


Figure 1 - MapReduce Example

The example uses the transaction log from a supermarket cash register to figure out how many times each item was sold. The system processes the data in several steps:

- **Input:** The system reads the log with all transactions, where each row represents a transaction.
- **Split:** The workload is split into individual transactions for more granular processing.
- **Map:** The mapping function extracts and counts each item from the transaction.
- **Combine:** The counts for the individual items are combined.
- **Reduce:** Aggregates all the items and creates the result.

This is only a small example with trivial processing, however, the presented programming model is able to connect thousands of cheap servers to create a huge pool of processing resources at a reasonable price. Since the first release of Hadoop's MapReduce many other data processing frameworks could build upon its achievements. The most notable frameworks are three projects maintained by the Apache Software foundation, i.e., Spark², Flink³ and Storm⁴.

The concepts for Spark have been first published in 2010 by Zaharia et al., while the first software release was in 2014 (Zaharia et al., 2010). The major difference to MapReduce is the introduction of the resilient distributed dataset (RDD), which enables in-memory processing and re-using of a data set for a huge performance increase across multiple parallel operations, e.g., iterative machine learning algorithms. Spark also allows data access via SQL for ad-hoc queries on big datasets, which caused significant delays on MapReduce.

Flink and Storm were both released in 2011. Even though they are smaller projects than Apache Spark, there are use cases where those frameworks are superior. The main differences between those projects is the computation model. Spark is designed to process a dataset at once, called batch-processing. Flink and Storm are stream-processing engines, which means they individually process a possibly infinite stream of incoming data tuples. Even though it is possible to use Spark for the

² Apache Spark Website: <https://spark.apache.org/> retrieved 08-11-2020

³ Apache Flink Website: <https://flink.apache.org/> retrieved 08-11-2020

⁴ Apache Storm Website: <https://storm.apache.org/> retrieved 08-11-2020

processing of micro-batches with just one data tuple, the native stream processing design of Flink and Storm lead to a lower latency in those use cases(Karimov et al., 2018).

The emergence of cloud computing and storage services, e.g., Amazon Web Services (AWS)⁵, Microsoft Azure⁶, Google Cloud Platform (GCP)⁷, further democratized the processing of big data. Companies that could not afford to buy a computing cluster outright can access the required resources on a cloud platform and just pay for the resources they use.

Further advances in virtualization of applications into containers led to the development of Kubernetes in 2014. Kubernetes is an open-source cluster management system based on Google's Borg that is maintained by the Cloud Native Computing Foundation⁸ (Verma et al., 2015). Kubernetes orchestrates data storage and processing locally or on any cloud platform with the highest grade of flexibility yet. Kubernetes is an important part of the CAAI architecture and will be explained in more detail later in chapter 6.3.2.

Manyika et al. created an overview of the possibilities for big data to create value for a company from the available data (Manyika et al., 2011):

- Big Data creates transparency as stakeholders can quickly access relevant data.
- It enables companies to set up experiments, e.g., experiments for process changes. Large amounts of data can be collected and analyzed to identify possible performance improvements.
- Big Data can be used to create a more detailed segmentation of customers to customize actions and prepare specific services.
- Analysis of Big Data can support human decision making by pointing to hidden correlations or hidden risks, e.g., risk or fraud analysis for insurance companies.
- Data can also enable new business models, products and services or can improve the existing ones. Data about how products and services are used can be used to develop and improve new versions of the product.

Auschitzky et al. developed a report with several use-cases to demonstrate the benefits of big data for manufacturing (Auschitzky et al., 2014). They found that “most companies collect vast troves of process data but typically use them only for tracking purposes, not as a basis for improving operations.” The readily available data from the shop-floor could be analyzed and brought significant improvements for the participating companies.

One of the biggest biopharmaceutical producers created a centralized data store to collect the process information for all products, process steps and used materials. The analysis created clusters of closely related production activities and further investigated each cluster. Statistical analysis found a set of nine parameters that were most influential for the production yield. The company was able to increase production yield by more than 50 % by targeted changes to those parameters.

A producer of specialty chemicals found that the levels of carbon dioxide in its processes had a significant influence on the overall yield and amount of waste. The advanced process controls helped reducing the waste of raw materials by 20 % and the energy costs by around 15 %.

⁵ Amazon Web Services Website: <https://aws.amazon.com/> retrieved 08-11-2020

⁶ Microsoft Azure Website: <https://azure.microsoft.com/> retrieved 08-11-2020

⁷ Google Cloud Platform Website: <https://cloud.google.com/> retrieved 08-11-2020

⁸ Cloud Native Computing Foundation Website: <https://www.cncf.io/> retrieved 08-11-2020

A precious metal mining company used Big Data to investigate the process control parameters and compare the results between different mines. The analysis showed that the levels of dissolved oxygen had a big impact on the extraction of precious metals. They adjusted the processes and increased the average yield by 3,7 % within three months without additional investments or major changes to the existing processes.

Modern manufacturing produces more and more data, and needs to ensure that it extracts and uses the value that lies in the vast amounts of data.

1.2 ARTIFICIAL INTELLIGENCE

Creating value out of all that available data is possible through AI. AI works by combining large amounts of data with fast, iterative processing and intelligent algorithms, allowing the software to learn automatically from patterns or features in the data (Bishop, 2006).

The term *artificial intelligence* was first used by the summer study group at Dartmouth College in 1956, organized by John McCarthy, where scientists met to discuss and clarify their ideas regarding thinking machines. The researchers in Dartmouth built on some of the even earlier works regarding intelligent machines of Turing, who tried to answer the question "can machines think?" in 1950 with a method that is known today as the *Turing Test*, and Bush, who proposed a system that enhances a human's knowledge and understanding (Bush, 1945; Turing, 1950). They founded the research discipline of *artificial intelligence* in 1955 on the assumption that "every aspect of learning or any other feature of intelligence can be so precisely described that a machine can be made to simulate it" (McCarthy et al., 2006; Moor et al., 2006).

Artificial intelligence itself is a broad research discipline, influenced by several other fields such as philosophy, mathematics, psychology, neuroscience, economics, and technical as well as theoretical computer science (S. J. Russell & Norvig, 2012). Over time several definitions have been created or adjusted and there are even proposals for alternative names for the research discipline itself:

- Barr and Feigenbaum propose the following definition: "Artificial Intelligence is the part of computer science concerned with designing intelligent computer systems, that is, systems that exhibit characteristics we associate with intelligence in human behavior – understanding language, learning, reasoning, solving problems, and so on" (Barr & Feigenbaum, 1981).
- Kaplan and Haenlein define artificial intelligence as "a system's ability to correctly interpret external data, to learn from such data, and to use those learnings to achieve specific goals and tasks through flexible adaptation" (Kaplan & Haenlein, 2019).
- Poole, Mackworth and Goebel suggest "computational intelligence" as alternative name for the field and define computational intelligence as "the study of the design of intelligent agents", where an intelligent agent is "a system that acts intelligently: What it does is appropriate for its circumstances and its goal, it is flexible to changing environments and changing goals, it learns from experience, and it makes appropriate choices given perceptual limitations and finite computation." They argue that artificial intelligence as the name for the field creates confusion as the method is to actually create intelligent systems but "the purpose is to understand how intelligent behavior is possible" (Poole et al., 1998).

Describing such an AI system today is mostly done using the agent metaphor as used by Poole, Mackworth and Goebel. Following this, Tecuci created Figure 1 which depicts the components of an

intelligent agent and the interaction with the environment, even though not all agents need to possess the same modules or use them to the same extend (Tecuci, 2012).

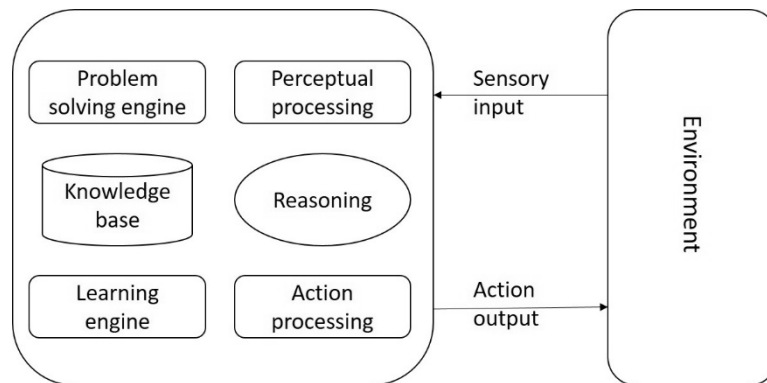


Figure 2 - The main modules of a knowledge-based agent

Poole, Mackworth and Goebel list the following inputs for such an agent to determine an output:

- Prior knowledge about the world
- Past experience that it can learn from
- Goals that it must try to achieve
- Observations about the current environment and itself

Those inputs enable the agent to learn with one of the types of learning explained in the following section.

1.2.1 Learning

For some very complex problems it is not feasible or very costly to create an algorithm in the traditional sense. The programmer would need to create a complete solution which foresees every possible event or obstacle and often the programs need to change over time to reflect modifications in the environment. The programmer may also not be capable to solve the problem himself. Therefore, it is often beneficial to use a learning algorithm, which derives knowledge from available data (Domingos, 2012).

The main types of learning are determined by the available feedback (S. Russell & Norvig, 2010):

- **Unsupervised learning:** The agent learns patterns in the input data, even without explicit feedback. The main purpose of unsupervised learning is clustering, where useful clusters are detected from the input examples. For example, a taxi driver agent could learn to distinguish "good traffic days" from "bad traffic days" without ever receiving labeled examples.
- **Reinforcement learning:** The agent receives a bonus or malus for its actions and tries to optimize the reward. After many iterations the agent learns which actions result in the best outcomes. For a taxi driver agent, it may be a hint for faulty behavior if it receives no tip after a trip.
- **Supervised learning:** The agent learns a function which maps input to output and tries to predict new instances. The taxi driver agent would study a dataset of trips, rated by the customer, to learn which is the best route for a given destination.

Even though each different approach to learning delivers impressive results in recent times, the beginning of AI was marked by very high expectations, which could not be immediately fulfilled and led to a long period called "AI Winter".

1.2.2 AI Winter and recent advances

The first results of AI in academia were very impressive. Newell and Simon developed a system that generates proves for most of the theorems in chapter two of Russel and Whiteheads "Principia Mathematica" (S. Russell & Norvig, 2010). Arthur Samuel developed a checker program which learned the game by playing against itself, human players, but also from book games to create a knowledgebase of roughly 53.000 positions. The program became "rather better-than average novice, but definitely not an expert" but demonstrated that "it will learn to play a better game of checkers than can be played by the person who wrote the program" (Samuel, 1959). Robinson created a "machine-oriented logic based on the resolution principle", which was, theoretically, able to prove any theorem using first-order logic (Robinson, 1965).

Unfortunately early research in AI was not able to live up to the very high expectations when transferred to real-world problems which led to the two so-called "AI Winters" from 1974-1980 and 1987-1993, where funding was cut after damning reports from the Automatic Language Processing Advisory Committee (ALPAC) of the U.S. government in 1966 and the Lighthill report for the British government in 1973 (Crevier, 1993).

During AI winter the research continued, largely isolated from each other, in a manifold of specialty fields. John A. Bullinaria composed an overview of relevant AI subfields (Bullinaria, 2005):

- **Neural Networks** – e.g. brain modelling, time series prediction, classification
- **Evolutionary Computation** – e.g. genetic algorithms, genetic programming
- **Vision** – e.g. object recognition, image understanding
- **Robotics** – e.g. intelligent control, autonomous exploration
- **Expert Systems** – e.g. decision support systems, teaching systems
- **Speech Processing** – e.g. speech recognition and production
- **Natural Language Processing** – e.g. machine translation
- **Planning** – e.g. scheduling, game playing
- **Machine Learning** – e.g. decision tree learning, version space learning

Enhanced algorithms contribute to the recent advances in all of those fields, but the huge amount of data and the capabilities to actually process it are even bigger drivers of the rapid progress since the early 2000s (Halevy et al., 2009). Banko and Brill worked on natural language processing (NLP) when they stated: "a mediocre algorithm trained on 100 million words performs better than the best available algorithm trained on 1 million words" (Banko & Brill, 2001). Hays and Efros, who worked on automatic scene completion in photographs, similarly conclude: "filling space in a picture through an algorithm did not work well when trained with 10.000 pictures, but achieved excellent results after training it with two million pictures" (Hays & Efros, 2007). The very distinct use cases display the possibility to solve a lot of problems through a learning algorithm if enough data is available.

In this regard state-of-the-art Deep Learning algorithms in image recognition and classification could improve the performance from a 26,2 % error rate in the late 1990s, to 15,3 % in 2012 and 2,25 % in 2017 using Convolutional Neural Networks (Ouaknine, 2018). A successful real-world application in healthcare is a software created by DeepMind, an Alphabet subsidiary, which identifies 50 eye diseases as accurately as a human doctor (Vincent, 2018).

The accelerated growth of available data and related progress also brings back hopes to develop artificial general intelligence.

1.2.3 Evolving from a very narrow AI towards artificial general intelligence

The first commercial applications for AI were expert systems, which captured human knowledge in a set of rules and were able to solve a very narrowly defined problem. One of the first commercial systems is R1, a system that helped configuring complex computer systems during the order process, introduced by the Digital Equipment Corporation in 1980 (McDermott, 1982). Those expert systems usually consist of three basic components: a knowledge database with facts and rules representing human knowledge and experience; an inference engine processing consultation and determining how inferences are being made; and an input/output interface for interactions with the user (Smith et al., 2006). By the end of the 1980s more than half of the Fortune 500 companies either developed or maintained expert systems (Enslow, 1989).

While the ultimate goal is reaching *artificial general intelligence*, an artificial intelligence that is capable of learning any intellectual task just like a human, most of today's AI systems are still trained to do a clearly defined task (Pennachin & Goertzel, 2007). The system that detects fraud cannot drive a car or give you legal advice. In fact, an AI system that detects health care fraud cannot accurately detect tax fraud or warranty claims fraud. However, recent advances enable the application of AIs on a broader field of problems, such as Watson, a famous AI showcase by IBM.

1.2.4 Famous AI Showcases and Competitions

IBMs Deep Blue Chess Computer defeated Gary Kasparow in 1997 with 3½–2½ in a match of six games. The parallel computer was able to generate up to 30 billion positions per move and "thought" 14 moves in advance (Campbell et al., 2002). While this was an impressive victory, critics attribute it to brute force not true intelligence. Brute force means the program was able to evaluate millions of possible future moves and their combination, much more than any human possibly could (Harmon, 2019; Hsu, 2002).

IBM developed Watson and showcased the system as a contestant in the TV game show "Jeopardy!" in 2010. In this game the players receive an "answer" and are required to find the correct "question" in a wide range of topics, such as history, literature, science, politics, film, art or pop culture. The game poses several challenges as the system needs to understand the natural language and the often subtle or ironic "answers", search the knowledge database for fitting "questions" and then determine which of the candidates is the best match. While Deep Blue was disassembled after the show match as it was not able to work in more generalized use-cases, Watson is now used as an AI platform and developers can use it to create new AI applications (Searle, 2011).

Deepmind, a subsidiary of Alphabet Inc., created AlphaGo to study the ancient board game Go and demonstrate their advances in AI. The game is played on a 19x19 board where the players alternate to place black or white "stones" and try to encompass and therefore control the most space on the board. As the players can place their stones anywhere on the board, Go has a huge number of possible developments (10^{170}) for every turn, which makes it a lot more complex than chess and experts thought of it as the holy grail of AI in games. At the end of 2015 AlphaGo was able to defeat Fan Hui, the European Go champion, five times in a match of five games, after studying historic Go games of professional players with the supervised learning approach and further self-play with reinforcement learning. After the first success it was possible to schedule another big show match in March 2016, just

four months later, against Lee Sedol. Nobody expected AlphaGo to win the match but the AI won with a convincing 4-1 and presented several inventive moves that led Lee Sedol to comment after an analysis of the games: "I thought AlphaGo was based on probability calculation and it was merely a machine. But when I saw this move I changed my mind. Surely AlphaGo is creative." Between the two show matches the AI had played millions of games against itself and was therefore able to improve rapidly. Those showcases always contain lessons that can be transferred to general applications for artificial intelligence.

The recent advances in AI and the huge amount of data that is generated every day allows rapid progress and solving more and more complex real-world problems.

1.3 INDUSTRY 4.0

I4.0 transforms manufacturing through integration of the digital into the physical world. In the future, smart factories collect more data than ever to empower AI in the next generation of production systems, also called cyber-physical production systems (CPPS). The main goals of I4.0 are improving production yield, reducing waste and additional productivity through increased automation (Spectral Engines, 2018).

The term I4.0 tries to highlight the importance of this progress, implicitly placing it as the successor of the three earlier industrial revolutions. Draht and Horch summarized those earlier stages as follows and Figure 2 from Spectral Engines shows an overview (Drath & Horch, 2014; Spectral Engines, 2018):

- **Industry 1.0:** Steam engines introduce mechanic help into fabric production in the 1780s. Production gets moved from workers' homes into central factories.
- **Industry 2.0:** The second industrial revolution stems from electricity and enabled mass production utilizing conveyor belts beginning in the 1870s. The most famous example for successful introduction of mass production and the resulting division of labor is the Ford Model T in the United States.
- **Industry 3.0:** In 1969 Modicon introduced the first programmable logic controller and therefore automation in manufacturing. The following advances in information technology brought immense additional benefits.

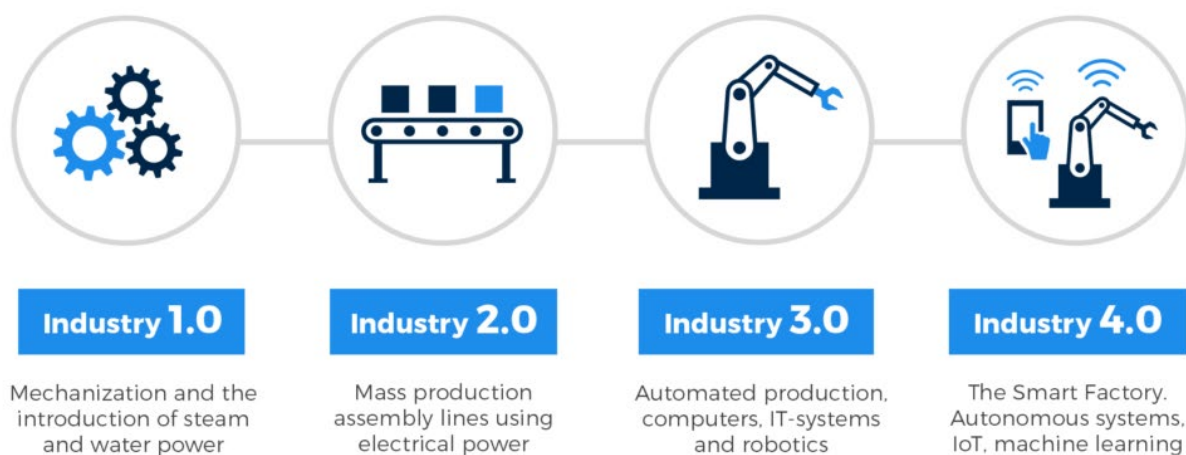


Figure 3 - The four industrial revolutions

As I4.0 describes the progress through new technologies and their cooperation, there is no clear definition and researchers characterize the key components instead.

Kagermann, Wahlster and Helbig identify the three main features of I4.0 as horizontal integration through value networks, end to end digital integration of engineering across the entire value chain, and vertical integration and networked manufacturing systems (Kagermann et al., 2013).

Gerbert et al. created a list of “9 Pillars of Industry 4.0” for their report published by the Boston Consulting Group in 2015 (Gerbert et al., 2015). Those pillars are: big data, autonomous robots, simulation, additive manufacturing, internet of things, cloud computing, augmented reality, horizontal and vertical integration and cyber security.

Herman, Pentek and Otto characterize the fourth industrial revolution by a paradigm shift from centrally controlled to decentralized production processes (Hermann et al., 2016). They identified the following key components of I4.0:

- IoT integration with manufacturing processes
- Cyber-physical System (CPS) through integration of computation and physical processes
- Smart factories through vertical integration of IoT and CPS in their operations

In summary I4.0 describes the advances in connectivity, enabling horizontal and vertical integration within a production plant but also between business partners, to create autonomous CPPS and make decisions utilizing the growing amount of available data and AI.

1.3.1 Cyber-physical production systems

"CPS are systems of collaborating computational entities which are in intensive connection with the surrounding physical world and its on-going processes, providing and using, at the same time, data-accessing and data-processing services available on the internet" (Monostori et al., 2016). If a CPS is used in a production setting it is called a CPPS (Bangemann et al., 2015).

They are at the core of I4.0 and “go beyond a single enterprise and cover the entire value chain and consist of an online network of sensors, machines, workpieces, and information technology (IT) systems” (Ramanathan, 2018). Due to increasingly available sensors and processing power it is possible to generate benefits from the additional data in several use cases, e.g., condition monitoring or optimization (Bunte, Fischbach, et al., 2019).

The U.S. national institute of standards and technology concludes that "the potential of CPS to change every aspect of life is enormous. Concepts such as autonomous car, robotic surgery, intelligent buildings, smart electric grid, smart manufacturing and implanted medical devices are just some of the practical examples that have already emerged" (National Institute of Standards and Technology, 2013).

1.3.2 Industry 4.0 example use cases

Practical industrial applications show the possibilities of I4.0 today. The German Mechanical Engineering Industry Association (VDMA) represents around 3.200 members, making it the largest industry association in Europe. They published a collection of over 50 use cases by their members which implement I4.0 as guideline and inspiration (VDMA, 2016). The different categories are presented with a selection of these use cases.

1.3.2.1 Product Customization

Customers wish for more and more individualized products and are willing to pay for it. Thus, mass customization has been identified as a competitive strategy by an increasing number of companies (Da Silveira et al., 2001). Industry 4.0 supports companies to create more flexible processes and produce those small batch sizes at a reasonably low cost.

The company Arburg experienced a growing demand for individualized plastic products, but those small batches could not be produced economically by their existing production line. Arburg combined injection molding, additive manufacturing and I4.0 technologies in a flexible and automated CPPS. The CPPS enables customization by producing the plastic product traditionally via injection molding and then 3D-printing the individual adjustments onto the workpiece. The result is the economic production of highly customized plastic products with a higher efficiency, process reliability and continuous quality assurance.

B&R produces automation technology, e.g. programmable logic controllers (PLC). Their customers demanded a greater variety, lower costs, better quality and a smaller batch size. B&R introduced fully automated production cells that produce more than 200 different PLC modules. Those CPPS are connected to the enterprise resource planning (ERP) systems and retrieve their instructions directly from the ERP order management. The outcome for B&R is an increase in productivity and profitability, integrated production tracking which improves the process transparency and product quality, and condition monitoring which enables predictive maintenance.

Kärcher produces floor cleaning machines in more than 40.000 variants. They introduced radio-frequency identification (RFID) chips in their production process, which carry the information for each individual order. The Kanban material flow system delivers the required parts to the workstation. The workstation retrieves the information for the individual product from the RFID chip and displays the assembly instructions for the worker. Finally the system automatically updates the order status when the production is finished and the unit passes the tests. Kärcher is now able to economically produce highly customized products through a more flexible production line with higher process transparency.

Industry 4.0 allows companies to follow their customer's demands and produce highly individualized products in an economic manner.

1.3.2.2 Predictive Maintenance

Real-time information about the operating status of machine components enables I4.0 to go from condition monitoring to forecasting the remaining time before a machine fails and planning the maintenance accordingly (Feldmann et al., 2017). Continuous measurements of important operating parameters help to optimize the maintenance timing and advance from "preventive maintenance", e.g. changing the brake pads after 100.000km, to a "condition-based maintenance", e.g. changing the pads based on the forecasted remaining kilometers before failure.

Aventics produces I4.0 valve systems, which generate diagnostic information and link all connected pneumatic valves, sensors and actuators to the higher level control system. The enhanced networking, even for devices which do not possess I4.0 capabilities by itself, enables condition monitoring and predictive maintenance for increased system availability.

Bachmann, who develop automation solutions for the machine building sector, take this monitoring one step further. They introduced the "Fleet management system" which collects monitoring data

from all their machines around the world to create global insights into their life cycle to enable predictive maintenance for their customers.

A mold identification module was introduced by the company Balluff. These molds for injection molding require regular maintenance and their RFID system makes the individual mold traceable. Their system can be fitted to every existing mold and tracks the number of produced items and process parameters. The data can be analyzed to conduct maintenance on the mold before it fails.

Festo, a company which produces industrial control and automation solutions, created a sensor module that is integrated in all new systems and tracks the energy consumption. It is able to detect savings potentials but also if a system is malfunctioning and therefore drawing more energy.

The Heidelberg Druckmaschinen AG produces industrial printing presses. They installed additional sensors in their machines to perform condition monitoring. The condition monitoring system detects approaching errors and prevents unplanned downtime. The performance data can also be used to optimize the production process. While their customers benefit from an increased throughput of up to 40 %, the Heidelberg Druckmaschinen AG also sees a growing service business and presents the predictive maintenance capabilities of their presses as unique selling point in comparison to their competitors.

All companies developed impressive solutions to prevent unplanned production downtime. While some solutions are built into new machines, others can be retro-fitted onto existing machines, which lowers the costs.

1.3.2.3 Product Life Cycle Management

The increased integration between systems enabled by I4.0 also leads to new possibilities for the Product Life Cycle Management. While today most departments within a company are not yet fully integrated, the I4.0 vision leads to a universal data-integration network, encompassing companies, suppliers and customers (F. Ferreira et al., 2016). Several companies presented projects that improve the integration through I4.0 technologies:

Bosch introduced a new system called “TraQ”, which stands for “Tracking and Quality”, to collect data on the shipping process of their products. The TraQ system collects data on the temperature and humidity of the environment and records shocks to the product. Supply chain partner can react to problems as they receive real-time alerts for the shipping conditions. Bosch benefits from continuous transport documentation, which improves the product quality and simplifies the complaints management.

Cognitas is specialized on documentation of manufacturing systems. Whereas manuals were created mainly for safety purposes in the past, now the operating companies need them for day-to-day operation as the systems become more and more complex. Cognitas created a system that generates a virtual image of a machine, called “digital twin” (Haag & Anderl, 2018). The image mirrors the actual machine state at any time and can be exchanged between companies. This leads to shorter maintenance and shutdown times while also increasing operation safety.

NanoFocus produces 3D microscope systems that can be integrated into the production line to control the process based on their measurements. This leads to a reliable and fully automated quality assurance, which does not depend on random samples, but instead evaluates each produced unit, which leads to higher quality products for their customers.

SAP integrated their manufacturing execution system (MES) with the production processes of ebm-papst in a collaboration of the two companies. The goal was to digitize the production and create a connection between the internal processes, suppliers and customers. This simplifies the order processing where the MES autonomously triggers production steps. The entire process can be managed in a central system where the production performance can be analyzed and visualized at any time. All the parts used in the product can be traced from the supplier to the customer as well as all processes between order and delivery.

The approaches are all very different, but highlight the possibilities of I4.0 technologies and the value that can be derived from the additional collected data.

1.3.2.4 Worker Assistance Systems

Industry 4.0 also impacts the relationship between humans and machines. While the manufacturing systems grow more complex, the product life cycles become shorter, making the operation and maintenance of the systems a challenge for the workers. Assistance systems support the workers with relevant and context-dependent information (Dhiman & Rocker, 2019).

Centigrade demonstrate their worker assistance system for intelligent production planning and logistics in a fully functional miniature production line. Workers can use a mobile device to see real-time information on available production resources and adjust the production plan at any time. Additionally cameras mounted on the ceiling of the warehouse create a live augmented-reality environment for the mobile devices. The worker can use his touch display to move objects on the screen and trigger actual operations in the warehouse, while gamification elements in the app try to motivate the worker. The assistance system leads to increases in efficiency while decreasing production risks at the same time.

The Centrum for Industrial IT in Lemgo also created a worker assistance system, but based it on projector technology. The assembly instructions are projected as 3D model onto the surface of the manual workstation, while lights show which parts should be used next. Thus, the versatile assembly system guides the worker safely and efficiently through the assembly process. Workers can switch the workstations more easily, which leads to a more flexible production.

The example use cases highlight how the growing complexity can be handled with short learning times through context-sensitive information and effective instructions.

1.3.2.5 Optimization

The optimization of existing processes can increase the energy efficiency, reduce the time spent to produce a product or shrink the delay between the production of different variants. I4.0 technologies and the collected data enable new use cases:

The company ebm-papst developed an intelligent control system for industrial refrigerators. Automatic transmission between the cooling units and the control system allows an optimization of the cooling times and best utilization of the cooling capacity. The resulting system saves 30 % of the energy spent and enables remote monitoring and maintenance.

KSB published an app for the smartphone to analyze potential energy savings in pump systems. The user simply specifies the type of the pump and motor and records the operation sounds for 20 seconds. The app transfers the data into the KSB cloud and analyzes the sounds. The result will tell the user if an optimization of the pump system is possible and advises for next steps.

Kuka presented the robot “iiwa”, which collects Kanban material boxes in the central warehouse rack. It uses a scanner for Quick Response (QR) codes to retrieve the target destination of the box and delivers the parts autonomously to the workstation. This saves the worker time and increases the assembly line efficiency.

All companies showed promising, innovative results and future developments will incorporate AI even deeper into CPPS.

1.3.2.6 Versatile Production Systems

The flexibility of traditional high-performance assembly lines was quite limited. A Versatile Production System (VPS) is a modular system with compatible interfaces, which allow different hardware configurations to enable adaptive production and the extension of the system with new modules.

Schunk build a collection of adaptable modules with freely programmable motors, grippers and sensors for a flexible design of production processes. Their VPS continuously monitors all process steps. Thus, process parameters can be corrected if necessary to ensure a high product quality and maximum productivity. The sensors detect deviations in the operation early, therefore prevent component failures and allow plannable and effective maintenance. The recorded data is also the basis to create additional services.

The demands for higher product customization requires new and more flexible production lines. The utilization of VPS supports companies by increasing their flexibility.

1.4 CONCLUSION

This chapter introduces the different fields big data, AI and I4.0. It gives an overview of the origins, developments and technologies of the respective fields, but also shows the possibilities for cooperation. The big data technologies allow the collection and processing of enormous amounts of data. AI learns from patterns in the data and the results improve drastically if it is able to train on a bigger data set. Additional sensors and processing power in manufacturing allow to use AI algorithms in I4.0 use cases. All the presented use cases show impressive first implementations of the new I4.0 technologies and highlight the future potential:

- Make already existing products smarter
- Create highly customizable products
- Offer the data generated from smart technologies as product or service
- Develop entirely new services or products
- Optimize production processes and performance

The next chapter evaluates how the technological progress translates into economic impact.

2 ECONOMIC IMPLICATIONS

This chapter examines the economic implications through the introduction of big data, AI and I4.0. Those new technologies influence individual companies but also impact the global economy. Therefore, both are examined before the resulting government programs are presented.

2.1 BRAND VALUE AND MARKET CAPITALIZATION

Successful utilization of new technologies and the vast amounts of available data is demonstrated by technology companies, such as Alphabet (former Google), Amazon or Apple.

Those companies dominate the rankings for the most valuable brands and the highest market capitalization. The value of a brand is defined by Aaker as “the short-run and long-run flow of profits that it can generate” (Aaker, 1991). The brand value is supported by different aspects, i.e. brand visibility, brand associations and customer loyalty. Brand visibility means that the brand is relevant to the customer. They are aware of the brand and think that it can credibly fulfill a certain need. Brand associations involve anything that created a positive or negative relationship with or feelings toward the brand and differentiate a company from its competitors and gives customers a reason to buy. Customer’s loyalty provides a flow of business for current and potential products from customers that believe in the value of the brand’s offerings and will not spend time evaluating options with lower prices.

In 2019 seven out of ten of the most valuable brands are technology companies, i.e., Alphabet, Amazon, Apple, Facebook, Microsoft from the U.S.A and Alibaba and Tencent from China, as shown in Figure 4 (Handley, 2019). In the same year Amazon overtook Alphabet (former Google) and Apple to become the most valuable brand at 315 billion dollars.

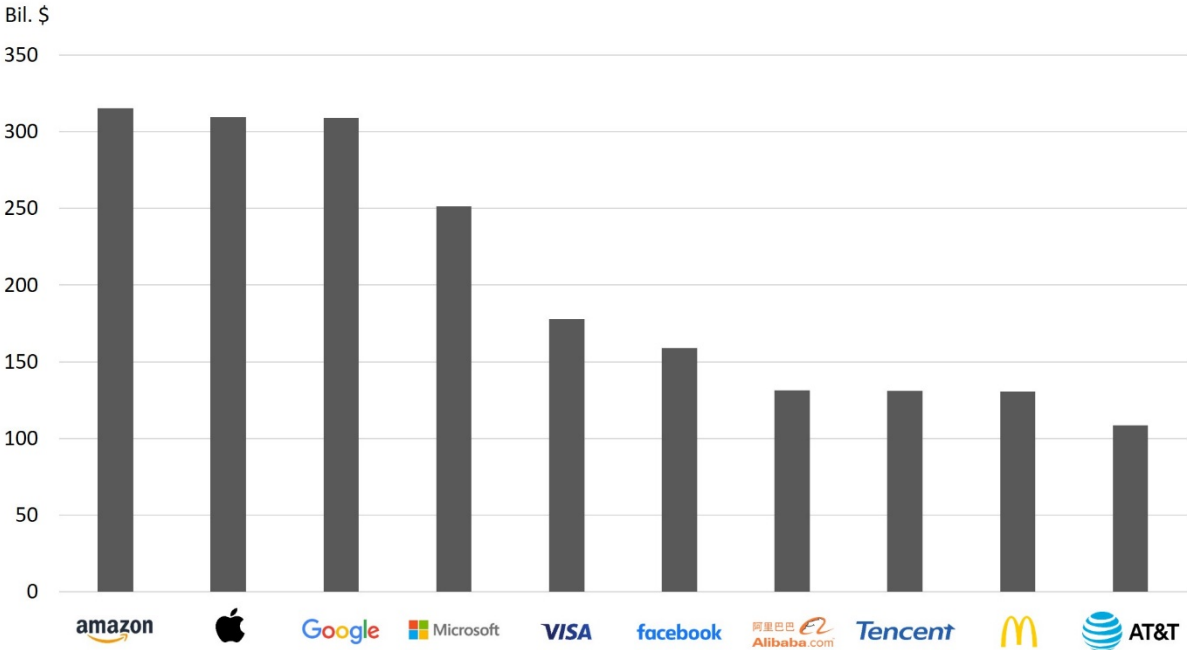


Figure 4 - Most valuable brands in 2019

Market capitalization⁹ represents the worth of a company as determined by the stock market. The stock market price of a company results from the supply and demand for shares of the company. The stock market price and thus market capitalization are significantly influenced by expectations on the future profitability of the company. Both metrics highlight the prospect of future earnings and development potentials. By now several technology companies also reached a market capitalization of more than a trillion \$US, only surpassed by the oil company Saudi Aramco, as shown in Figure 5 (Winck, 2020). Big Tech, the five stock group comprised of Alphabet, Amazon, Apple, Facebook, and Microsoft, passed a \$5 trillion market valuation for the first time at the beginning of 2020.

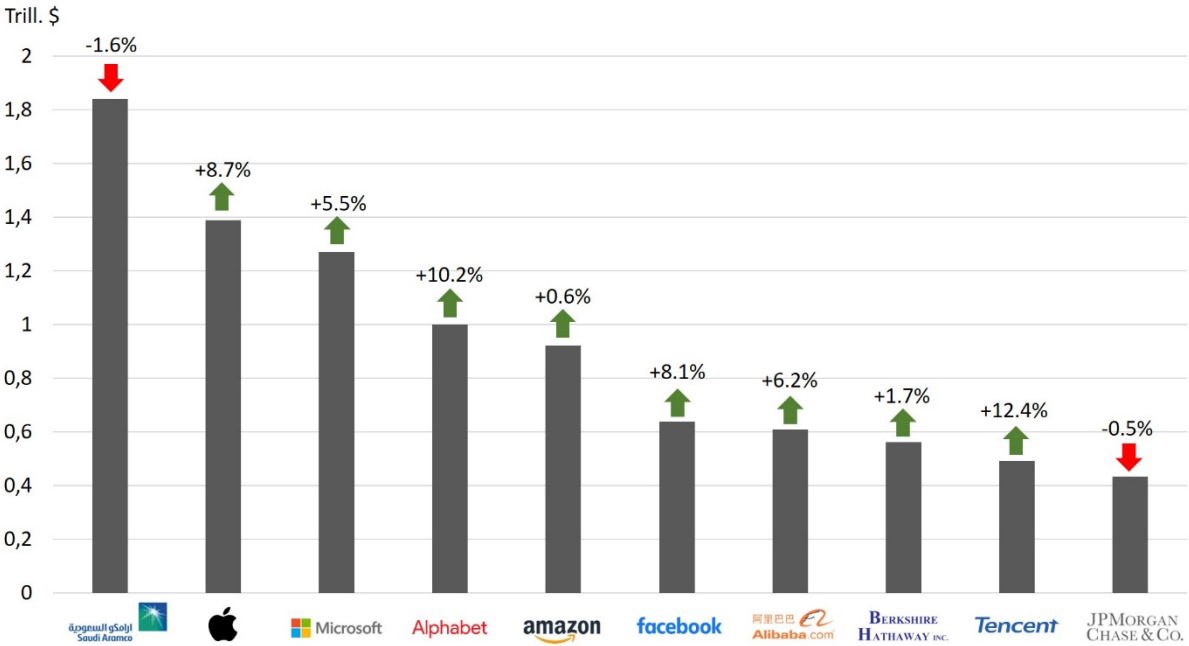


Figure 5 - Market capitalization with year-to-date difference

Amazon is a prime example for companies that successfully utilize new technologies and the collected data. What started 1994 in a garage as a digital shop for books quickly turned into the biggest e-commerce platform in the world. But since then Amazon constantly grew and evolved their business model and also became market-leader for cloud computing, digital streaming and artificial intelligence among many others.

Successful digital companies are less “anchored” to individual categories or regions. Instead they offer a wide variety of products and services around the world (Handley, 2019). Often the more services they offer the more value they can extract from the collected customer data. However, the potential of the new technologies influences not only individual companies, but entire economies.

⁹ Market capitalization is defined as total number of outstanding shares multiplied by the Dollar value of a single share, see <https://www.investopedia.com/terms/m/marketcapitalization.asp>

2.2 GDP AND PRODUCTIVITY

Advanced economies experience a slowing economic growth, measured as increase of the gross domestic product¹⁰ (GDP). Main factors for this development are slowing growth of the labor force, additional investments and also productivity¹¹ (Arsov & Watson, 2019). Furman investigated the productivity growth rates of advanced economies and compared data from 1996 to 2006 and 2006 to 2016 (Furman, 2017). In the later time period 36 out of 37 economies had slower productivity growth, averaging a 1.0 % annual growth compared to an earlier average annual growth of 2.7 %.

New technologies and changing business models are required to further increase the productivity in advanced economies. Each of the following technologies allows companies to increase their productivity and adapt or reposition their business model. But only the combination of these technologies shows the real potential. This impact potential on the global economy has been investigated by several researchers in recent years.

The successful utilization of big data and AI leads to a predicted increase of roughly 15% in global economic output by 2030 (Manyika et al., 2018; Rao & Verweij, 2017). Manyika, Chui and Joshi created a simulation that forecasts effects on GDP and labor markets until 2030 which add around 16 % or 13 trillion dollars compared with today, as shown in the figure below.




Figure 6 - Simulated economic impact of AI, based on (Manyika et al., 2018)

The simulation derives this number by combination of an estimated 26 % GDP increase and the related costs, e.g., technology introduction and labor displacement. They calculate an average annual increase of 1.2 % of GDP which results in modest effects in the first years of implementation but substantial effects later on. The introduction of steam engines in the 1800s in comparison resulted in an annual

¹⁰ Gross domestic product measures the monetary value of goods and services produced in a country. The growth of GDP is widely used as reference point for the health of an economy and compares the output of given time periods, e.g., per quarter or per year, defined in <https://www.imf.org/external/pubs/ft/fandd/basics/gdp.htm>

¹¹ Productivity is commonly defined as a ratio of a volume measure of output to a volume measure of input use, see <http://www.oecd.org/dataoecd/59/29/2352458.pdf>

productivity increase of 0.3 % and the widespread introduction of IT into business during the 2000s resulted in 0.6 % more productivity per year.

Rao and Verweij built a large scale model of the global economy based on the Global Trade Analysis Project database (Rao & Verweij, 2017). The database consists of trade data for 140 countries and their transactions in 57 economic sectors. Their research predicts a 14 % increase in global GDP until 2030 and also assesses the gains for specific countries, as shown in their figure below.

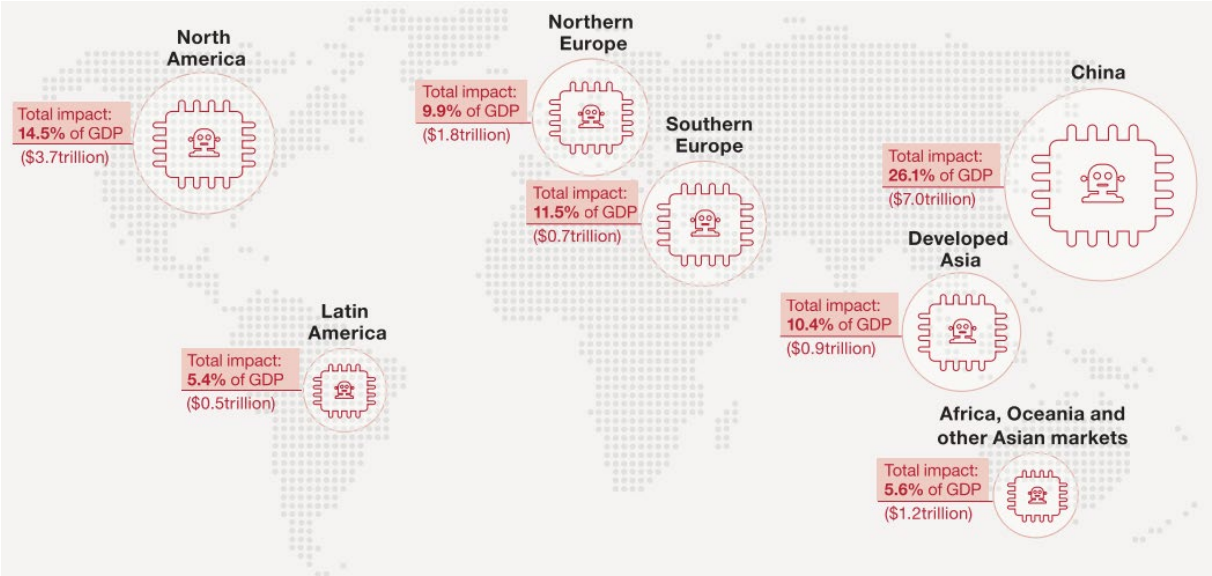


Figure 7 - Economic impact of AI introduction per region

Rao and Verweij anticipate a potential 14 % GDP boost for the U.S.A and an increase of up to 26 % for China. The predicted increase for the latter is exceptionally high because introducing new production technologies impacts the large proportion of Chinese GDP generated through the manufacturing sector. The researchers expect that America quickly realizes the AI benefits while China needs more time to accumulate the technology and knowledge, but will progress faster by 2030. This assessment seems reasonable as most of the former AI breakthroughs were US-led and many countries currently develop AI initiatives to become competitive (Manyika et al., 2018).

McKinsey Global Institute published a discussion paper in which they mapped traditional analytics and AI techniques to over 400 use cases in companies and organizations to assess the practical applications and the economic potential of advanced AI techniques across industries (Chui et al., 2018). Among their use cases 69 % could be improved through AI beyond traditional analytic techniques, 16 % could be improved only with AI and in 15 % AI could not improve the results further than with traditional analytics, often because of data limitations. Their research shows that almost three-quarters of the impact from advanced analytics is tied to use cases requiring millions of labeled data examples. This means that organizations will have to adopt and implement strategies that enable them to collect, integrate, and process data at scale. They find the greatest potential value impact from using AI in marketing and sales but also in operational functions, including supply-chain management and manufacturing. The potential impact for manufacturing stems from predictive maintenance,

optimization of production yield, inventory and parts, logistics and warehouse management, but also improved procurement, sales and demand forecasts.

Both studies found manufacturing to be among the sectors with the largest productivity gains from AI and identified three areas with especially high potential: enhanced monitoring and auto-correction of manufacturing processes, supply chain and production optimization and on-demand production. All of those areas are addressed in I4.0.

Sunken costs for sensors as well as IT processing capabilities enable the widespread collection of process- and production-related data and therefore intelligent manufacturing. I4.0 profits from the progress made by IoT, which is a related technological development with a focus on a global connection of physical and virtual assets. The advances of IoT required compact, durable and energy-efficient hardware and resulted in cheap sensors and single-board computers that can also be used in a Smart Factory. The worldwide economic value of IoT is estimated to reach \$1.9 trillion by 2020 while the prices of processors are decreasing to roughly one dollar (Slama et al., 2015). Similarly, the costs for sensors fell nearly 200 % between 2004 and 2018 to an average of 0.44\$, which makes intelligent manufacturing affordable not only for big multinational companies but also for smaller manufacturers (Microsoft, 2018).

Using those connected devices in manufacturing systems, e.g. to monitor the movements of workpieces on the factory floor or to get real-time updates on equipment status, is expected to improve operational efficiency by 2.5 % - 5 % (Manyika et al., 2013). As the number of connected machine-to-machine devices increased by 300 % from 2008 to 2013, Manyika, Chui and Bughin see the potential for “almost every (80-100%)” manufacturing facility to use connected devices until 2025 which would lead to an economic impact of 900 billion up to 2.3 trillion dollars in cost savings per year.

Bauernhansl created a detailed list of possible cost savings for different areas of production through introduction of I4.0 (Bauernhansl, 2017):

| Type of cost | Total value |
|---------------------|----------------|
| Inventory costs | -30 % to -40 % |
| Manufacturing costs | -10 % to -20 % |
| Logistical costs | -10 % to -20 % |
| Complexity costs | -60 % to -70 % |
| Quality costs | -10 % to -20 % |
| Maintenance costs | -20 % to -30 % |

Table 1 - Possible cost reductions through I4.0

The inventory costs can be reduced as the optimal amount of inventory can be adjusted based on real-time information and a lower safety stock is needed. Manufacturing costs decrease through optimized usage of equipment, measured as overall equipment effectiveness, and personnel. Additional automation will cut the logistical costs as more and more autonomous transport systems move goods to their destination. Bauernhansl predicts the highest cost reductions through reduced complexity. To lower complexity the product portfolio needs to be examined to find the most profitable products and services with a subsequent analysis to identify growth potentials and synergies between similar products. Real-time information about the production process allows timely adjustments and increases the overall quality of the product, which decreases costs related to faulty and rejected

workpieces. Predictive maintenance is another I4.0 field with a significant savings potential for every manufacturer. Maintenance can be planned and prioritized according to the actual status of the equipment and the predicted remaining time before failure. This prevents unforeseen downtimes, reduces the necessary inventory of replacement parts and allows the technicians to be better prepared for any maintenance task. Companies need to upgrade their machinery and computational infrastructure to benefit from those saving potentials and business consulting firm Frost&Sullivan estimate that the supplier ecosystem of I4.0 solutions alone will reach 420 billion euros in value by 2020 (Frost&Sullivan, 2015).

I4.0 not only represents cost savings or productivity increases for the companies, but also enables and accelerates the emergence of new products and services. Lichtblau et al., conducted a study with German manufacturing companies in 2015 and the majority (60%) stated that I4.0 will allow them to increase their revenues by implementing new business models (Lichtblau et al., 2015). Correspondingly Wischmann, Wangler and Botthoff expect additional annual sales related to I4.0 of 20 to 30 billion euros for German manufacturing companies (Wischmann et al., 2015).

2.3 GOVERNMENT INITIATIVES

Governments worldwide start national research initiatives to introduce big data, AI and I4.0 and gain a competitive advantage in the global economy as Manyika et al. suggest that “a sizable portion of innovation gains come as a result of competition that shifts market share from non-adopters to front-runners” (Manyika et al., 2018).

Most of the former AI breakthroughs were achieved in the US and in this context President Trump signed the Executive Order “Maintaining American Leadership in Artificial Intelligence” in February 2019 with a budget totaling \$973.5 million in 2020, and planned budget raises for the following years, across a range of agencies, programs, and initiatives (The White House, 2019). The White House also recently published a detailed report on the initiative where they outline the key areas for AI research (The White House, 2020):

- Long-term investments in AI research
- Developing effective methods for human-AI collaboration
- Addressing the ethical, legal and societal implications of AI
- Ensuring the safety and security of AI systems
- Developing shared public datasets and environments for AI training and testing
- Measuring and evaluating AI technologies through standards and benchmarks
- Better understanding the national AI R&D workforce needs
- Expanding public-private partnerships to accelerate advances in AI

While the document remarks that the United States continue their efforts to implement the OECD principles (OECD, 2020) for trustworthy, innovative AI that respects human rights and democratic values, they also express the urgency to remove obstructions of innovation:

“Not using AI technologies because of perceived or potential harms, however, could be just as problematic, depriving individuals — or the Nation — of the significant benefits that AI technologies could bring.”

The document also highlights the importance of an AI-ready workforce, which requires an enhanced focus on the STEM education talent pipeline, as well as technical apprenticeships, reskilling, and lifelong learning programs to better match Americans' skills with the needs of industry.

China published the "new generation AI" guideline and created a "national team" consisting of the three internet giants Alibaba, Baidu and Tencent to rapidly progress the AI industry with the goal to create a \$150 billion domestic market until 2020 and becoming a major center for AI innovation by 2030 (The State Council of the People's Republic of China, 2018). European Union members also announced to collaborate on AI technologies, in addition to existing national initiatives, to ensure that Europe stays competitive and to establish solutions for the legal and ethical challenges together (European Commission, 2020a, 2020d).

The German AI strategy was publicly presented at the 2018 Digital Summit (BMW, 2018). The initiative is funded by the Federal Government of Germany with around €3 billion for the period 2019-2025 and aims to make Germany a "leading center for AI, especially through pursuit of a speedy and comprehensive transfer of research findings to applications". They particularly talk about the transfer of AI research results to SMEs as the small and medium sized companies represent a huge potential for optimization. The strategy also highlights the necessity of a human-centric AI design that allows the development of workers' skills and talents, enabling self-determination, providing security and protecting human health. This includes large-scale qualifications initiatives with attention for lifelong learning expand and upgrade AI-related skills of the workforce.

All developments should be done for the citizens to minimize risks associated with the changes and create a "responsible use of AI in a way that serves the good of society, doing so together with science, business, the public sector and civil society [...] based on European values, such as the inviolability of human dignity, respect for privacy and the principle of equality". The basis for the responsible use are the recommendations from the Data Ethics Commission and the General Data Protection Regulation (GDPR) as reliable legal framework and further examinations whether the regulatory framework needs adjustments. They conclude that successful introduction of AI requires transparency of the processes used and an understanding of how it works and why it is beneficial.

The EU Commission summarized the goals in their recent report (European Commission, 2020b):

- Increasing and consolidating Germany's future competitiveness by making Germany and Europe a leading center in AI
- Guaranteeing a responsible development and deployment of AI which serves the good of society
- Integrating AI in society in ethical, legal, cultural and institutional terms in the context of a broad societal dialogue and active political measures

The Spanish Government presented a series of priorities and recommendations in 2019 that will be implemented within the new Spanish Strategy for Science, Technology and Innovation ("EECTI") until 2028 and financed through the Science, Technology and Innovation Stares Plans("PECTI") to mobilize the synergies between public and private sectors and avoid bias or discrimination in AI (Gobierno de Espana, 2019). The strategy sets out to create an framework for AI development and the measurement of its impact on business and society. The Spanish Government emphasizes the need for training and professionalization in the field of AI to ensure the transfer of knowledge and profits to the public.

The EU Commission recently published a report on the Spanish AI strategy and highlighted (European Commission, 2020c) several promising initiatives:

- Founding of a National Data Institute for a better data governance and improved quality of public services
- Development of a digital data ecosystem consisting of high-quality databases with open data access
- Creation of a Spanish Committee on Research Ethics to draft an ethics code for AI and to define ethical guidelines for a fair and sustainable use and development of AI.

Whole countries, as well as individual companies, need to assess the new technologies and re-think their digitalization strategy to maximize the economic benefits. Research by Rao and Verweij on the economic impact of AI found three driving factors (Rao & Verweij, 2017). In the near-term the biggest drivers are productivity gains from increased automation and businesses augmenting their existing labor force with AI technology. But in the long term they predict even bigger profits from consumer demand for more personalized or higher quality products and services.

Additional research projects and government programs have been initiated with a focus on I4.0. In the United States I4.0 is included in the broader terms Industrial Internet of Things (IIoT), Smart Production or the Industrial Internet. For the US administration I4.0 is a lower priority and not necessarily regarded as a key technology for future competitiveness therefore the progress is largely driven by private sector initiatives such as the Industrial Internet Consortium (IIC) (Kagermann et al., 2016). The IIC was founded in 2014 by General Electric, AT&T, Cisco and IBM and could quickly attract more than 200 national and international members. The IIC aims to progress common architectures, interoperability and open standards and established a cooperation initiative with the German Plattform Industrie 4.0 to achieve these goals together.

In 2015 Chinese government announced its Made in China 2025 strategy to transform the current mass production economy into a high-tech economy (Kagermann et al., 2016). China realized a GDP of 11.4 billion US Dollars in 2015 and is the second largest economy in the world after the United States. However, the industrial production sector represents a much bigger part of this GDP than for any other country. Manufacturing industry accounted for around 43 % of GDP in 2014, compared to 21 % for the United States and 31 % in Germany (Urmersbach, 2019). China aims to fully modernize its manufacturing industry to improve production efficiency and quality via technological advances and identified automation and digitalization of Chinese industry as a key enabler.

Germany launched the national initiative “Industrie 4.0” in 2011 based on the government’s “High Tech 2020” strategy (European Commission, 2017a). The Ministry of Education and Research (BMBF) and the Ministry for Economic Affairs and Energy (BMWi) aim to increase digitization and interconnection of products, value chains and business models to establish digital manufacturing over a 10-15 year period. The German Academy of Science and Engineering (acatech) collaborated with experts and published the research agenda and implementation strategy for I4.0 in 2013 (Kagermann et al., 2013). Since then the German government guaranteed more than €470 million in support for I4.0 research and implementations (BMBF, 2018). The government emphasizes that industry financing is essential for the success of the initiative. Thus, companies partially finance the research projects they participate in, based on their company size. The stated overall goal is to “become a leading supplier of smart manufacturing technologies as well as to develop new lead markets for CPS technologies and products”.

Spain announced “Industria Connectada 4.0” in 2014 “aiming at digitizing and enhancing the competitiveness of Spain’s industrial sector” (European Commission, 2017b), an initiative driven by the General Secretary of Industry and SME. The government defines three main goals:

- Increase the added value and employment in the industrial sector
- Develop and encourage the use of a “Spanish model” for the industry of the future
- Develop digital solutions and competitive levers to promote the Spanish industry and increase exports

The initiative is defined as public-private initiative and the government supported companies with €170 million in loans and direct aid in 2016. The initiative emphasizes a focus on implementation of digital transformation projects in SMEs and micro enterprises through financial support and personalized guidance in this process.

2.4 CONCLUSION

The combination of big data, AI and I4.0 provides significant potential for every manufacturing company. Cheaper sensors and processing power make intelligent manufacturing affordable for companies of any size. Implementing those technologies in high-wage countries’ manufacturing companies maintains their technological advantages and helps to compete in a global market (Manhart, 2013). The ability to produce many distinct products in small batch sizes or with a batch size of one has to improve to remain competitive (D. T. Matt & Rauch, 2013; Spath et al., 2013). Potential customers expect product customization, which requires a more flexible and connected supply chain. (Baum, 2013). Companies need highly adaptable manufacturing systems to quickly respond to customers’ wishes (Zawadzki & Zywicki, 2016). Initially, those expectations lead to a more complicated and challenging situation for manufacturing companies. However, the expected profits are also higher than from any industrial revolution before, and research suggests that a considerable amount of innovation gains stem from the redistribution of market shares from non-adopters to front-runners (Manyika et al., 2018). The introduction of I4.0 and AI in manufacturing companies makes high-wage countries competitive on a global market. (Kagermann et al., 2013).

However, despite the enormous economic benefits the introduction of AI and I4.0 technology also present ethical challenges that need to be discussed. The next section gives an overview of the ethical challenges.

3 ETHICAL IMPLICATIONS

AI and I4.0 will be introduced because the economic benefits are massive and countries must stay competitive on a global market. The development introduces new challenges for policy makers, who have an interest in the higher labor productivity, economic growth and prosperity, but need to ensure responsible use of AI technology and the related data (Chui et al., 2018).

Policy makers as well as company owners need to involve the factory workers and take their needs into consideration to gain their support instead of creating modern Luddites. The history of the Luddites is detailed in an article by Richard Conniff which was published in the Smithsonian Magazine (Conniff, 2014). The movement of the Luddites was born at the beginning of the 19th century when many British working class families suffered from widespread unemployment. The first industrial revolution introduced machines into manufacturing and the textile industry was the first to introduce modern production methods. Protests on 11.03.1811 in Nottingham, a textile manufacturing center, for more work and better wages were broken up by the British military which triggered the destruction of textile machinery in that same night. What began as solitary raids eventually became a wave of attacks that spread across a 70 mile radius. The government feared a national movement and sent thousands of soldiers to defend factories. Bloody battles between Luddites, factory owners and the military continued until 1816 and many workers were sent to the gallows or into exile. But the Luddites did not randomly attack factories. They were also not opposed to technology or not able to use it. Many were highly skilled machine operators in the textile industry. They attacked manufacturers who used machines in “fraudulent and deceitful manner” and fought for “machines that made high-quality goods [...] run by workers who had gone through an apprenticeship and got paid decent wages” (Binfield, 2004).

Policy makers also need to check if their historical tools and frameworks are still adequate to handle the consequences of the new technologies for the common worker as their jobs change or are replaced by machines. Some of the tools that are regularly discussed in this context are the universal basic income (De Wispelaere & Stirton, 2004), wage supplements (Phelps, 1997) or a guaranteed employment (Mosler, 1998). The degree of disruption caused by AI and I4.0 on the labor markets determines which tool is most suitable or if entirely new frameworks are necessary.

3.1 INFLUENCE OF AI AND I4.0 ON THE GLOBAL LABOR MARKET

Researchers portray different scenarios how those new technologies will influence the global job market. Frey and Osborne were the first to investigate the effects and predicted for the American workers that 47 % of their occupations can be automated (Frey & Osborne, 2013).

In 2015 Bonin, Gregory and Zieran transferred the studies from Frey and Osborne into a German context and conclude with a similar estimate of 42 % (Bonin et al., 2015). However, they clarify that they do not think the result can be equated with a loss of 42 % of all jobs, because it is not the occupation as a whole that can be automated but individual tasks. Therefore, they present an alternative calculation and conclude that 12 % of all jobs are highly probable to be automated in the future.

Bughin et al. analyzed 46 economies for possible impact and their research suggests between zero and one-third of work activities could be displaced by 2030, with an average of 15 %. They predict a greater effect on advanced economies, as higher wages imply more economic incentive for automation.

According to Bughin et al. between 75 million and 375 million workers (3 % - 14 % of the global workforce) will experience a major transition towards new or other occupations. This would match or even exceed the scale of historical shifts where technological advances revolutionized agriculture or manufacturing (Bughin et al., 2017).

Bessen later argues that the new technologies will have a positive effect on overall employment if they improve the productivity in markets with unmet demand. So while employment may decline in manufacturing, the spill-over effect into other industries could lead to a positive job creation, similar to other types of automation (Bessen, 2018).

Robotics is such a related type of automation and may be a reasonable indicator for the consequences of AI and I4.0 for the labor market. Robots also see increasing use in industry as the prices decrease to just a fraction of their original costs (Graetz & Michaels, 2018), while usage becomes more flexible and functionality grows. Subsequently worldwide robot shipments increase about 150 % between 2010 and 2016 (Furman & Robert Seamans, 2018). The council of economic advisers analyzed the use of robots in several countries in 2012 and identified the automotive sector as most prominent. They calculated the number of robots per 10.000 workers in the automotive sector and found 1.091 robots in the United States, while Japan utilized most robots (1.563) and Germany employed slightly more robots than the United States (1.133). In the same time the average for all other industries in the United States was just 76 robots per 10.000 workers (Council of Economic Advisers, 2016). Dauth, Findeisen, Südekum, and Wößner researched the influence of robots on the labor market and combine German labor market data from 1994 to 2014 with IFR robot shipment data. They identified changes in the composition of work activities in the manufacturing industry due to the introduction of robotics. Those changes lead to higher wages for engineers and management but lower wages for low- and medium-skilled workers and less new jobs for young labor market entrants. This development increases social inequality as highly skilled workers and capital owners are strongly favored by the introduction of robots. They also estimate that each additional industrial robot, while increasing the overall work productivity, leads to the loss of two manufacturing jobs. However, they found no negative impact on the overall employment level because enough new jobs are created in the service industry to offset and in some cases over-compensate for the negative employment effect in manufacturing. (Dauth et al., 2017).

The creation of new jobs is not limited to the introduction of robotics in manufacturing, but a general trend. Research by Lin shows, that between 1980 and 2000 about 4 % - 9 % of the United States workforce were employed in jobs that did simply not exist 10 to 15 years earlier (J. Lin, 2011).

Thus, workers need the formal education and further training to adjust to the changes in work activities and be able to use the new technologies. Governments and companies must, in their own interest, do everything possible to prepare workers for life-long learning and provide the training needed to keep up with the changing requirements

3.2 THE NEW DIGITAL DIVIDE

The introduction of big data, AI and I4.0 creates new jobs with new requirements for the employees. The technology companies Google, Microsoft and Facebook each hired several thousand workers to evaluate new AI systems and ensure reliable services (The Economist, 2017). The growing ecosystem of I4.0 solutions similarly needs engineers, who design the intelligent components to build a modern CPPS and operate the resulting systems.

Bach et al. conducted a cluster analysis on European economies and found that the northern European countries endured less negative effects of automation, as the jobs in these countries tended to be more complex and therefore harder to automate. A high level of digital development should prevent negative consequences of the new technologies both at the country and the company level. However, they also state that a low level of digital development reinforces the negative impact and creates a “digital divide” (Bach et al., 2020).

While the “digital divide”, first introduced in 1990, originally meant access to the internet, the “new digital divide” evolved towards the knowledge about and the ability to use algorithms, AI and the machines that incorporate those. This divide favors prosperous countries, which can afford extensive research initiatives, big multinational companies, that have the resources to fund expensive R&D departments, and individuals with a higher education.

One of the groups at a major disadvantage are workers with a low level of education as mostly routine tasks will be replaced. This is no new phenomenon but a long term trend, represented by a decline of the male labor force participation rate in the United States since the 1950s. This rate has fallen from 98 % in the 1950s to 89 % in 2016 (Council of Economic Advisers, 2016). The Council of Economic Advisers found a strong relationship between occupations and skills that can be automated and income or education. They used the characterizations from Frey and Osborne and found that 83 % of jobs with an income of less than 20 dollars per hour could be automated while jobs with an income of more than 40 dollars per hour only had a 4 % probability of automation. Furman also found that most of the men who do not participate on the labor market have a high school degree or less, which suggests that they find it difficult to learn additional skills and transition to a new occupation.

Learning new skills and adjusting to changes in the workspace will be crucial as technology may replace specific tasks rather than entire jobs – leaving substantial room for human employment in jobs that will be changed by workers having a new tool at their disposal (Furman & Robert Seamans, 2018).

Instead of manually operating machines, the factory worker of the future will largely specify and monitor production strategies for CPPS (Gorecky et al., 2014). However, those new work activities need to engage and motivate the workers to create a positive and sustainably productive work environment.

3.3 QUALITY OF WORK AND WORK ALIENATION

The quality of work life as well as the mental and physical health of the employees and the influences on productivity is an active research field for many years. In 1967 Seeman concluded that work, which provides for little self-control, meaning, and intrinsic satisfaction leads to reduced motivation in the work process and to various forms of withdrawal (Seeman, 1967). Cummings and Manring investigated the relationship between those feelings of alienation and their impact on work-related behavior (Cummings & Manring, 1977).

Their study conducted with forging workers was based on Shephard's dimensions of alienation (Shepard, 1972):

- **Powerlessness** – the feeling that an individual is dominated and controlled by other people or a technical system of production
- **Meaninglessness** – the own work does not seem to progress the goals of the organization
- **Normlessness** – the feeling that career progress is not based on performance or competence
- **Instrumental work orientation** – the employee does not feel fulfilled or satisfied by the job
- **Self-evaluative involvement** – the extent to which a job is central to a person's identity

They were able to prove a significant relationship between powerlessness, normlessness, and meaninglessness towards the worker's motivation and work performance. The less power the job holder feels in his job or over his career progress, the lower his effort and performance. The more the worker feels that his role is not related to the goals of the organization, the lower his effort and performance.

Mullarkey et al. investigated the additional stress for the machine operators that results from the introduction of advanced manufacturing technologies (Mullarkey et al., 1997). They define two technology characteristics that influence the worker's experience, i.e. technological uncertainty and technological abstractness. Technological uncertainty describes the extent of operational problems the advanced manufacturing technologies will have. Technological abstractness refers to the difficulties to understand and interpret the machines behavior in case of operational problems. When the information in more complex technology becomes less visible to the worker, the abstractness increases. They conducted the study in a manufacturing enterprise in England and found that the impact of uncertainty and abstractness relates to the pacing of the machines and, thus, the workload of the machine operator. Mullarkey et al. created Figure 8, which describes the relationship. When the machine pacing was high but the process was really stable, so the uncertainty was low, the workers enjoyed the continuous work and reported the development of a "flow" state. When the machine pacing was equally high, but the operation had frequent interruptions, the uncertainty increased and the workers felt that they fell behind schedule while solving the current problem. A slow pace of the machine alters the work environment completely. When the process is very stable and there is no uncertainty or abstractness the employees experience their work as monotonous. They cannot use their skills and are bored. However, if the machine pace is slow but the operation is complicated the workers feel every problem as a relief from boredom and as a challenge to use their skills. The results show the relationship between technology characteristics and the work environment and present the impact on the worker. Creating work situations with "Traction" and "Challenges" leads to satisfied workers, while the employees suffer from "Distraction" and "Passivity".

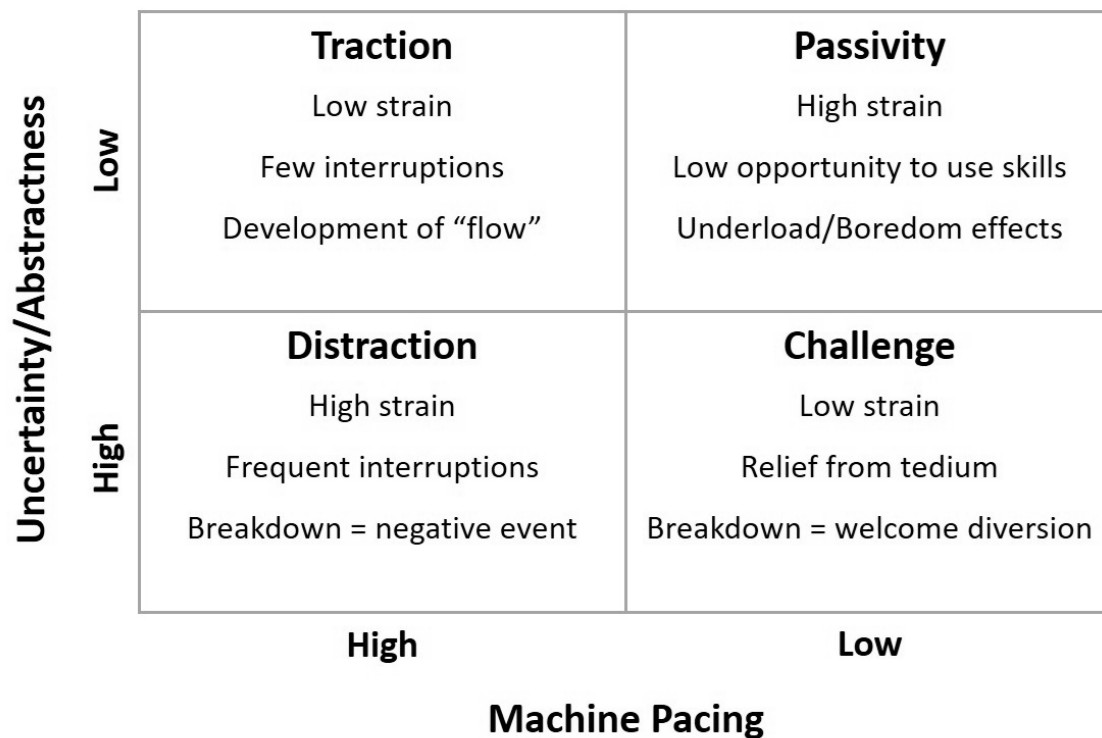


Figure 8 - Model of the relationship between machine pacing and technological characteristics

Research on organizational identification, job satisfaction, job involvement, job effort, job performance, the intention to quit, organizational turnover and personal alienation followed shortly after. Danna and Griffin reviewed the existing literature on health and well-being in the workplace and compiled evidence on the consequences of low levels of health and well-being (Danna & Griffin, 1999). They found absenteeism, reduced productivity and efficiency, reduced product and service quality and high costs for health insurance and medical expenses.

Sanderson and Andrews reviewed prevalence studies of mental health conditions and found depression among the most prevalent disorders in working populations around the world (Sanderson & Andrews, 2006). Reducing the rate of depression among workers is critical as mental health disorders are affecting millions of individuals worldwide and “associated with increased health care costs as well as productivity losses in the form of absenteeism, short-term disability absences and reduced on-the-job productivity – known as presenteeism” (Burton et al., 2008). Sobocki et al. researched the impact of depression on the Swedish economy and found that the costs associated with sick absence and early retirements had doubled since 1997 (Sobocki et al., 2007).

Thus, the promotion of health and the prevention of health problems is beneficial for a company. A good work environment should induce low levels of stress and reward the workers (Semmer, 2006). Semmer found several factors that create stress, such as high pressure, social conflicts, difficulties to accomplish tasks or lack of resources, such as control, social support and recognition. He groups possible changes to the organization of work into three categories: task characteristics, work conditions and social conditions. The tasks should be demanding, but not too complex, instead of boring and repetitive. Work conditions describe ergonomic issues, e.g. office furniture, noise reduction, worktime and workload. Social conditions refer to social relationships, conflicts with others, or lack of social support. While Semmer found that it is possible to achieve positive health-related

results, he notes that a combination of interventions on the individual and organizational level with support from company leadership are necessary.

Kobayashi et al. introduced the Mental Health Action Checklist (MHACL) to improve the mental health of workers in a Japanese manufacturing enterprise with a participatory approach as previous studies had shown favorable effects, e.g. reduction of stress, depression and employee absence (Kobayashi et al., 2008). They found MHACL helpful as it guided the supervisors as well as the workers and improved quantitative job overload, supervisor support, coworker support, and depression in the departments with a high participation rate (50% or higher) compared to the reference group.

Greenhaus et al. on the other hand found first evidence that employee happiness improves employee productivity (Greenhaus et al., 1987). Later Eskildsen and Dahlgaard investigated the relationship between quality of work, employee involvement and employee satisfaction (Eskildsen & Dahlgaard, 2000). They concluded that more involvement leads to more satisfied employees, which are highly motivated, demonstrate good morale and are more productive.

The rise of the “knowledge economy” established employee satisfaction and employee loyalty as critical issues because knowledge is a highly mobile resource (Matzler et al., 2003). Previous studies found that genuine trust in the workplace has several positive effects, e.g. higher employee satisfaction and commitment, more knowledge sharing, higher individual and group performance. Little trust in comparison was disturbing and upsetting for the employees (Dirks & Ferrin, 2002). Thus, trust should have a strong and direct impact on the employee satisfaction and ultimately the employee loyalty. Matzler and Renzl investigated the relationship of interpersonal trust, employee satisfaction and employee loyalty (Matzler & Renzl, 2006). For their research they used the definition of trust by Whitener et al. with three facets that comprise trust (Whitener et al., 1998):

- first, trust in another person reflects a person’s expectation or belief that the exchange partner will act benevolently;
- secondly, trust involves the willingness to be vulnerable and risk that the other person may not fulfill the expectations;
- and thirdly, trust involves a certain level of dependency which means that a person is affected by the actions of others.

Their studies could confirm a strong link between trust, employee satisfaction and employee loyalty and the impact of trust between colleagues was even greater than the positive impact of trust towards management.

The research presented in this section highlights the various influences on the well-being of an employee and the impact on work performance. A low quality of work and the employee’s alienation from work lead to productivity losses in the form of absenteeism, short-term disability absences and reduced on-the-job productivity. Creating a positive work environment does not only increase the mental and physical health of the workers, it also increases the productivity and fosters employee loyalty. A famous quote by Anne M. Mulcahy, former CEO of Xerox, summarizes the findings:

“Employees are a company’s greatest asset—they’re your competitive advantage. You want to attract and retain the best; provide them with encouragement, stimulus, and make them feel that they are an integral part of the company’s mission.”

Building trust not only between employees, but also between humans and technology becomes more and more important as these systems perform a wider range of tasks and the interaction with intelligent systems increases.

3.4 ACCOUNTABILITY AND TRUST IN NEW TECHNOLOGIES

Using the continuously growing amount of data and artificial intelligence in the I4.0 environment also transfers the related ethical issues into the field of advanced manufacturing.

Accountability

I4.0 connects machines within a plant, the suppliers and customers, and the resulting CPPS become increasingly complex to satisfy the new requirements. CPPS are safety critical systems, but who is responsible when the embedded AI or the whole CPPS fails at its task? Many different entities build the parts for the machine, develop the artificial intelligence and operate the finished system. Thus, the “many hands problem” also applies to AI in CPPS as “tracing the ethical responsibility and decision making of each stakeholder [...] poses a major challenge” (Siau & Wang, 2020; Wang & Siau, 2019). Pieters and Cleeff defined the “many hands problem” as follows (Pieters & Cleeff, 2009):

“In a complex chain of events or systems, many people will have had a share in an action that leads to undesirable consequences. As such many people will also have had the opportunity to prevent these consequences, and therefore no-one can be held responsible.”

The definition highlights the accountability problems in complex manufacturing systems. Earlier research by Helen Nissenbaum focuses on the challenges of accountability in software development (Nissenbaum, 1996). She also identified the “many hands problem” as the first of four issues. The second issue is the inevitability of bugs in software and the attitude of software developers towards thorough testing. The third and fourth issue stem from blaming the machine instead of the involved humans and simply refusing accountability.

The situation becomes even more complicated when artificial intelligence is involved, as the software developers build the base model but then the system continues to learn on its own and may come to the wrong conclusions or form an inherent bias (Wang & Siau, 2019).

Bias and discrimination in artificial intelligence

Artificial intelligence algorithms should be free from bias, as they make more and more important decisions and impact lives on a daily basis. Unfortunately, this is not the case and the outcomes often discriminate against specific groups, e.g., based on gender (Koolen & Van Cranenburgh, 2017; Larson, 2017) or race (Eliot, 2020). The bias stems from the way artificial intelligence learns from the data. In the most simple form, the AI searches for patterns in the data and tries to learn through generalization. If the bias is represented in the data the AI will “inherit” it. For example, a company wants to support the human resources department through artificial intelligence in the hiring process. They collect the historic data from all earlier applicants and label the data with “accepted” or “rejected”. Later the company discovers that mostly white men are hired for management positions, regardless of their qualifications. The inherent bias stems from the data, where historically mostly white men were hired

for management positions. Therefore, the AI learned, that the best criteria to determine good candidates may be the gender or race.

Siau and Wang concluded in this regard that “AI agents are only as good as the data human put into them” (Siau & Wang, 2020).

Unfortunately, the more complex the machine learning algorithm, the harder it is to explain the factors that lead to a decision and detect bias in a model. Paul de Laat therefore proposes a fifth issue for the age of big data and artificial intelligence: the opacity of algorithmic decision making (De Laat, 2018).

The opacity of algorithms can be mitigated with additional transparency to ensure the interpretability of models. This challenge has rekindled the debate about "explainable AI" and Adadi and Berrada laid out four different motivations for explainable AI (Adadi & Berrada, 2018):

- **Explain to justify**
The AI provides justifications about particular decisions to avoid bias and discrimination. The “right to explanation” introduced with the GDPR across the EU in 2018 makes this behavior mandatory.
- **Explain to control**
Insights into the decision making process make errors easier to detect and correct.
- **Explain to improve**
Continuous improvements to artificial intelligence models through better understanding.
- **Explain to discover**
AI can surpass human capabilities and generate new knowledge (see AlphaGo in 1.2.4). Explainable models can transfer the additional knowledge back to humans.

Adadi and Berrada conclude that explainability helps to justify decisions made by an AI, helps to improve the models and gain a better understanding about the problem itself. Thus, it supports the creation of more trustworthy AI systems. Building trust towards AI and its implementation in CPPS is especially important as this leads to technology acceptance, a critical factor for the introduction of I4.0

Building trust in technology

Understanding the results of an algorithm is a first step to build trust between a human and the machine that incorporates the algorithm. However, building trust in technology is a complicated process.

The definition of trust by Whitener (see 3.3), which describes the human characteristics, is extended with two additional dimensions, i.e. environment characteristics and technology characteristics, to describe the building of trust in technology (Hengstler et al., 2016; Oleson et al., 2011; Schaefer et al., 2016). Environmental characteristics concern the task that should be fulfilled and the team with its individual members and their attitude towards technology, which may or may not favor the integration of technology. Technology characteristics are described either as performance-based or attribute-based. A system with lots of errors will not be trusted as the human believes it is unreliable. Attribute-based characteristics concern the design of the system, e.g., a robot that resembles a human encourages trust.

Several studies concluded that trust is not gained in a single interaction, but build up in a series of interactions (Li et al., 2008; McKnight et al., 1998). Siau and Wang split the trust-building process in two phases, i.e., initial trust formation and continuous trust formation (Siau & Wang, 2018):

- Initial trust is based on the appearance of the technology, e.g. a humanoid robot, or good reviews from other users.
- Continuous trust formation favors AI applications that are easy to use and reliable. Implementing a collaborating system that is able to explain its actions will foster additional trust. In contrast, a system that does not respect the user's privacy or threatens his job will not gain trust.

Companies need to address the challenges of job security and user privacy to gather trust in the technology from their employees. They need to make sure that the AI's goals match the human goals as a pre-condition for trust. Even more, the workers would like to know the "state of mind" of their virtual co-worker and favor explicable algorithms. Whether employees can build trust in the technology also depends heavily on how the collaboration at work is designed.

3.5 DESIGNING THE RELATIONSHIP BETWEEN HUMANS AND MACHINES

The application of big data, AI and I4.0 will change the way we work, as intelligent systems remove a lot of repetitive activities and require human collaboration for complex tasks as well as creative thinking. Buhr summarizes the different scenarios for the possible relationship between humans and machines in the I4.0 age (Buhr, 2015):

- The "Automation scenario" assumes that technology will take control of the production process while the value of human labor declines. The CPPS decides which menial tasks can be carried out by human workers.
- In the "Hybrid scenario" workers and machines work cooperatively and complement each other, as both possess unique strengths for the production process.
- The "Specialization scenario" advocates for skilled labor as the deciding agent, using the CPPS as a supporting tool, to improve the overall process.

The scenarios depict very different relationships between workers and technology. The first scenario is outright dystopian from the workers point of view, but also very unlikely. As Furman and Seamans point out, AI is not developing fast enough to have dramatic effects on employment (Furman & Robert Seamans, 2018). AI is still far from an "Artificial General Intelligence", an intelligence that is far superior to humans in every aspect and capable to solve any use case through extensive learning. Right now, AI can only be trained for very specific tasks and workers' uncodeable practical knowledge cannot be replaced by smart technologies, which is of absolute importance to ensure stable production (Pfeiffer & Suphan, 2015). Therefore, the second or third scenario seem much more probable and designing a fulfilling workplace for autonomous employees is crucial.

Employees help companies realize their digital transformation and are the ones most affected by the changes of the digital workplace. Their direct working environment is altered, requiring them to acquire new skills and qualifications. Abel, Hirsch-Kreinsen and Steglich explain the worker's doubts not only with their fear of job losses, but also the technological changes being digital and no longer immediately comprehensible to the individuals which results in insecurities and skepticism (Abel et al., 2019). Kagermann, Wahlster, and Helbig similarly report a "growing tension between the virtual world and the world of workers' own experience. This tension could result in workers experiencing a loss of control and a sense of alienation from their work as a result of the progressive dematerialization and virtualization of business and work processes" (Kagermann et al., 2013). They also agree that through extensive human-machine interactions, the work content, and processes, as well as the working environment, will be radically transformed and thus also the worker's job and competence profiles.

Other researchers perceive the introduction of I4.0 not only as a challenge but also as a chance to improve the work environment. Schröder suggests a human-centered implementation with active involvement of the affected workers to create a positive relationship between humans and machines and thus increase the technology acceptance (Schröder, 2016). Gorecky et al. similarly propose to implement learning systems, which “dynamically detect and adapt to the context of the support situation and the worker’s actions” (Gorecky et al., 2014). In conclusion, for a successful introduction of I4.0 in any company, the employee technology acceptance got identified as one of the most critical aspects. Arias-Oliva et al. confirmed in their study, that innovativeness is a key variable in new technology acceptance (Arias-Oliva et al., 2020). Other factors are the perceived usefulness and ease of use. All these aspects must be explained to the employees, which requires communication and transparency as “acceptance is a fragile construct, which needs constant cultivation to convert employee resistance into acceptance or even support” (Abel et al., 2019).

3.6 THE ETHICAL INDUSTRY 4.0 INTRODUCTION PROCESS

The preceding chapter highlights the various ethical implications and challenges the introduction of new technology has on global economies, but also on individual companies. This section summarizes the steps and precautions a company can take to successfully introduce I4.0 in their organization.

The introduction of I4.0 is a major effort, that affects all parts of a company, e.g., strategy and leadership, logistics and supply chain management, customer relationship management, production planning and production processes, IT, technology and innovation management, and human resources (HR). Thus, a holistic approach is required, that includes all layers of management. Transparent communication includes the employees in the introduction process and relieves the worry of job losses.

The HR department develops the strategic approach towards the worker’s effective employment in the I4.0 environment. The promotion of health and the prevention of health problems is beneficial for a company as worker productivity is highly depending on the mental health and the “Mental Health Action Checklist”, introduced in section 3.3, is one of the tools for HR professionals to support and improve the worker’s mental health.

Learning new skills and adjusting to changes in the workspace will be crucial to overcome the negative effects of increasing automation, as shown in section 3.2, and for a successful I4.0 introduction. A systematic approach is therefore needed to teach employees the necessary skills. Furthermore, as the required competences grow more and more complex, it becomes obligatory to fit the qualification program to the individual worker. Hecklau et al. developed a competence model for human resource management in the I4.0 environment that enables the assessment of individual skills and competencies, as shown in Figure 9 (Hecklau et al., 2016). They derived four competence clusters, i.e., technical competencies, personal competencies, methodological competencies and social competencies, and the associated required skills. This makes it possible to adapt the training measures to the skills of the respective employee and thus adjust and optimize the training process for each worker.

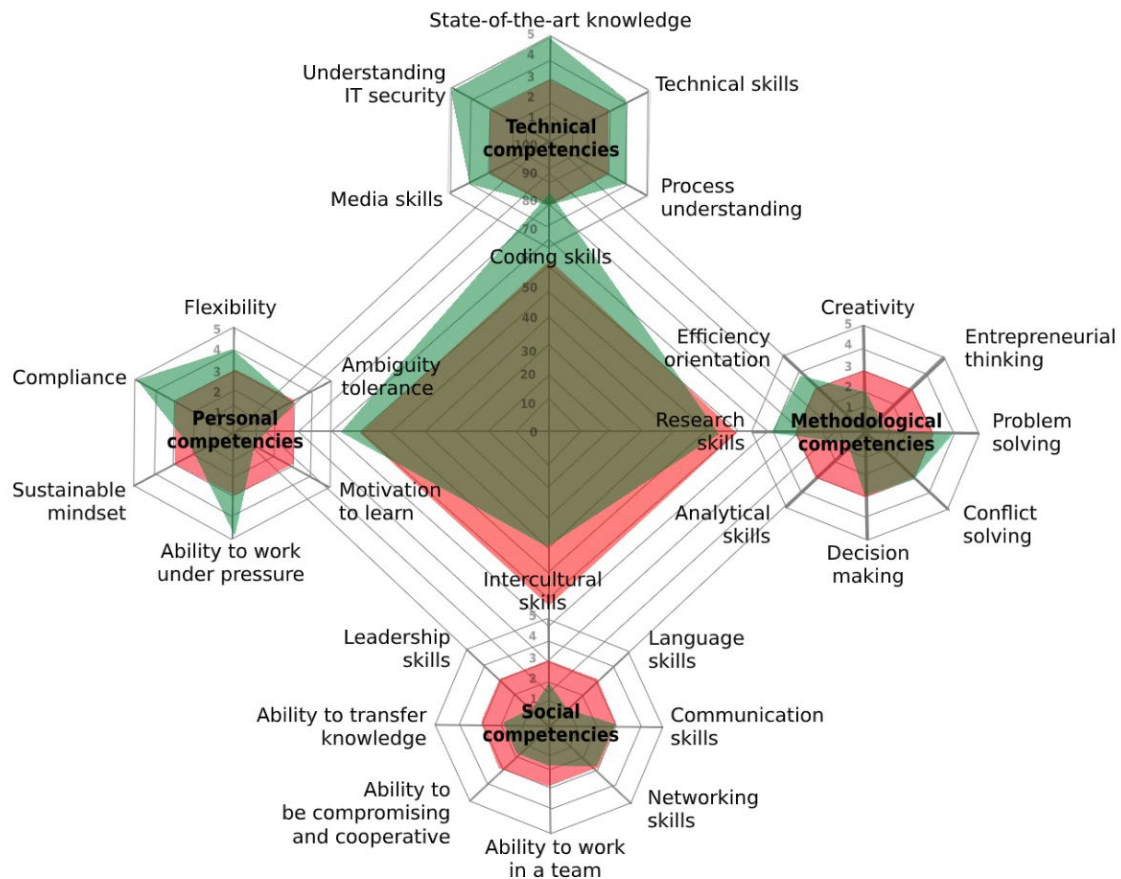


Figure 9 - Employee competence model

The ethical I4.0 introduction process also needs to consider the quality of work to create a good work environment that induces low levels of stress and instead rewards the workers. Low levels of stress does not mean, that the workers should not be challenged, as shown in section 3.3, but instead that workers need to have adequate time to solve arising problems and grow from the challenge. Thus, the workers should be included in the design of processes to benefit from their process knowledge and create processes that support the workers instead of overburdening them.

The introduction of I4.0 and the associated increase in automation also means that the cooperation between man and machine must be re-thought. Designing the relationship in a human-centric way, as shown in section 3.5, will prevent work alienation and instead ensure motivated workers that perceive the new technology as support and not as a threat.

The fruitful cooperation between human workers and machines requires humans to trust the machines. Therefore, the new technologies need to display predictable behavior and accountability as shown in 3.4. Building trust in the “new colleague” will further increase the technology acceptance rate.

Considering these aspects in the introduction process will help to overcome the challenges of increased automation and instead create a work environment that assists the workers and emphasizes the strengths of workers and machines.

3.7 CONCLUSION

This chapter discusses the various ethical implications new technologies such as big data, AI and I4.0 introduce. Policy makers and companies need to work together to create an environment that favors life-long learning and ensures worker's employability. Furthermore, it is in everyone's interest to introduce the new technologies in a human-centric way to provide fulfilling jobs to healthy and productive workers. A collaboration between humans and machines can emphasize individual strengths. However, for a fruitful cooperation it is necessary to build the trust in new technologies. Transparency is one way to create trust in technology and explainable AI helps to build trust through additional information on the decision making process. Even more, some companies release their software as open source software, which should be the preferred way, as noted by Paul de Laat (De Laat, 2018). One of the advantages of open source software is that everyone can analyze the source code which allows to find and debate features with an ethical impact. The interdisciplinary work is important as engineers who build new CPPS and the software engineers who create AI embedded in the manufacturing systems do not have the training to spot and solve ethical issues (Wang & Siau, 2019). Many eyes with different perspectives on the ethical issues can help to transform those into ethical solutions as "understanding and addressing ethical and moral issues related to AI is still in the infancy stage" (Siau & Wang, 2020). Reliable behavior is another important factor to build trust in new technology, or as Grodzinsky et al. put it "we prefer boringly predictable AIs" (Grodzinsky et al., 2010).

Raising awareness on ethical issues and creating trustworthy solutions will advance the technology acceptance, which is critical for I4.0 introduction. The ethical I4.0 introduction process, as shown in 3.6, supports companies in the introduction of the new technologies in a way that includes the workers and expands their competencies. The next section introduces the "Industry 4.0 Questionnaire for SMEs" that was conducted to get insight into the status of the I4.0 implementation in small and medium-sized companies and their attitude towards the new technologies.

4 INDUSTRY 4.0 QUESTIONNAIRE FOR SMEs

The introduction of I4.0 in companies is mostly shaped by big multinational enterprises. However, SMEs are just as important for economies around the globe. In the 28 European countries two thirds of employees in the non-financial sector are employed by SMEs, with 29.1 % in micro enterprises with less than 10 persons employed and 20.2 % in small enterprises with less than 49 persons employed and 17.1 % in medium enterprises with less than 250 employees. All three sizes of SMEs are contributing nearly equally to value added in the EU-28 with 20.3 % for micro enterprises, 17.6 % for small enterprises and 18.5 % for medium-sized enterprises (Airaksinen et al., 2015). In 2015 they employed 91 million people in total and generated 3.934 € billion of value added (Eurostat, 2018). Therefore “SMEs are the backbone of the European and many other economies” (*Future Image Industry 4.0*, 2012; Kraemer-Eis & Passaris, 2015). This is also especially true for manufacturing where European SMEs provide around 45 % of the value added and around 59 % of employment (Vidosav, 2014).

Even though SMEs are an important factor to economies, the I4.0 methods are developed mainly in larger enterprises and have to be adapted to the specific requirements of SMEs (Rauch et al., 2018). Currently the spread of I4.0 depends on company size and large companies are more likely to deploy relevant I4.0 technologies than SMEs (Schröder, 2016). Several researchers investigated the reasons for this observation. In 2017 Decker conducted a case study to evaluate the I4.0 readiness of Danish SMEs from the metal processing sector with the result, that SMEs at this time were not sure if or how they should introduce I4.0 in their companies (Decker, 2016). Wuest et al. confirmed the struggle of SMEs to adopt I4.0 in a study conducted with manufacturing SMEs in West Virginia in 2018 (Wuest et al., 2018). Both studies seem to support the claim made by Lutz Sommer in an article from 2015 that, “actually most of SMEs are not prepared to implement I4.0 concepts” (Sommer, 2015). Further research suggests various challenges for SMEs in I4.0 introduction:

- Different prerequisites regarding the integration of their product plants in higher level IT systems, which is much more advanced in bigger companies (Lichtblau et al., 2015).
- Using a self-assessment tool is not easy as I4.0 concepts are still too little known (Rauch et al., 2018)
- SMEs often lack resources to evaluate new technologies and their business uses. Thus it is hard for them to develop an appropriate strategy including a cost-benefit analysis (Schröder, 2016).

Those challenges need to be verified and addressed because “successful implementation of an industrial revolution I4.0 has to take place not only in large enterprises but in particular in SMEs” (Sommer, 2015).

Therefore, "Industry 4.0 will play a significant role in transforming traditional companies into smart factories with the help of IoT and CPS" (Erboz, 2018). However, in 2018, just 14 % of 1.600 executives, who participated in a study carried out by Deloitte (Deloitte, 2018), believed that their organization is prepared for I4.0 and able to profit from this new potential. Samaranayake, Ramanathan and Laosirihongthong conducted a study on the successful I4.0 introduction in companies and the most important factors they found were either process or technology related (Samaranayake et al., 2018). To benefit from the introduction, the production process needs to be stable, but also flexible enough to allow for an optimization. Equally important is the technology knowledge of humans and the ability to manage big data.

The questionnaire examines the technology acceptance by employees and the overall status of I4.0 introduction in SMEs. Several studies and questionnaires investigate the introduction of I4.0 and AI in manufacturing companies, but are mostly targeted towards big multi-national companies. Therefore, they provide the basis for an extended questionnaire that focuses on SMEs and analyses the employee's feelings in more detail.

4.1 INDUSTRY 4.0 SELF-ASSESSMENT

Companies can use various "maturity" or "readiness" models for self-assessment of their current I4.0 capabilities and progress towards successful I4.0 introduction. Schumacher, Erol and Sihn created an overview of existing models in 2016 and found that many models lack details regarding the development process or assessment methodology. They highlight the "Industry 4.0 Readiness Model" (Lichtblau et al., 2015) as "scientifically well-grounded and its structure and results explained in transparent manners" but the model contains just a single question for the employee dimension. In this question, they assess if the workers have the required skills to accomplish their future tasks (Impuls Stiftung, 2015). Schumacher, Erol and Sihn propose their own "Industry 4.0 maturity model", which also asks for the openness of employees towards new technologies but this approach could be further improved by using the technology acceptance model.

4.2 TECHNOLOGY ACCEPTANCE MODEL

The "Technology Acceptance Model" (TAM), developed by Fred D. Davis in 1989, is used to acquire additional insights into the factors that influence adoption of new technology (Davis, 1989). The two main variables are "perceived usefulness" and "perceived ease-of-use". Im, Kim, and Han extended the TAM by introducing "perceived risk" as an additional variable that negatively affects adoption (Im, Kim, & Han, 2007). Those factors influence the "attitude towards usage" and finally also the "behavioral intention to use" and the questionnaire contains items to investigate every component. To this day TAM is one of the most popular models to assess user acceptance of new technologies and was successfully used to evaluate the adoption of related technologies, e.g., smartphones and wearables (Chang, Lee, Ji, 2016; Roy, 2017).

4.3 METHODOLOGY

The questionnaire is conducted in Germany and Spain with a focus on SMEs. The German manufacturing companies are all members of the "Innovation Hub Oberbergischer Kreis", a regional association that focuses on the exchange of I4.0 knowledge and possible applications. The Spanish companies belong to the industrial service sector, working on optimization, logistics and manufacturing technologies for their national and international clients.

The questionnaires are sent to the executives, who are responsible for the introduction of I4.0 and AI in their companies. The complete questionnaire is available in the Appendix in English and online in Spanish, English and German for further reference¹².

¹² Jan Strohschein, Github Repository: <https://github.com/janstrohschein/Industry-4.0-readiness-for-SMEs-Questionnaire> Retrieved at 18-10-2020

The presented challenges for I4.0 introduction of SMEs influence the hypotheses (H1-H7) for the questionnaire. Previous research suggested that smaller companies require more assistance for the I4.0 introduction as the integration between production systems and higher level IT systems is often limited and resources are scarce. Different aspects of assistance are assessed in H1-H3 to get a better overview of the required support. Thus, H1 investigates if smaller SMEs need more assistance for evaluation of I4.0 technologies than bigger companies. As smaller SMEs already use less technologies and are less familiar with the integration of higher level IT systems it is expected that it will be harder for them to evaluate additional new technologies. Introducing new technologies comes with certain costs which leads to the second hypothesis as smaller SMEs may also require more assistance to assess the costs and benefits of I4.0 technologies. The resulting benefits depend on the company and the production processes. The costs for the acquisition and introduction of the new technologies are also based on the existing technologies and the current IT environment. The realistic assessment of costs and benefits will be harder if no experiences with the introduction and integration of higher level IT systems and intelligent manufacturing technologies exist. It is therefore expected that smaller companies need more assistance for the cost and benefit analysis of the I4.0 introduction. The third investigated assistance aspect is the formulation of a company-wide I4.0 strategy. Smaller companies that formerly did not use or need a process for innovation and technology management could find it harder to develop a I4.0 strategy across departments. Thus, it is expected that those companies require more assistance than bigger companies.

The development of I4.0 technologies is currently mostly shaped by big multinational companies, as those can afford to operate research and development departments. Smaller companies may profit from the experiences of bigger partners as long as best practices and standards are not widely available. The exchange with a bigger partner may also demonstrate the usefulness of I4.0 technologies and, subsequently, raise the technology acceptance. Thus, H4 examines if the collaboration between SMEs and bigger partners supports the I4.0 introduction for SMEs and increases the technology acceptance rate. Both is expected to be verified.

The hypotheses H5 and H6 investigate the effect of the participating companies' motivation to introduce I4.0. Motivation can have internal and external origins. Internal motivation describes behavior that is rewarding in itself, e.g., a company strives to improve the quality of their products because they want to deliver the best products. External motivation depicts behavior that is influenced through a third party, e.g., a company introduces I4.0 because they fear that their competitors will otherwise get an advantage. Thus, H5 examines if employees from a company with an internal motivation to introduce I4.0 have a higher technology acceptance rate. As internal motivation is more rewarding than external motivation it is expected that the employees have a higher technology acceptance. H6 assesses if companies with an internal motivation to introduce I4.0 have higher expectations regarding the increase in productivity through the new technologies. It is expected that the companies with an internal motivation have higher expectations as they introduce I4.0 because they recognize the potential for their company.

The last hypothesis investigates the differences regarding the I4.0 introduction between Spanish and German participants. The German government launched the national program "Industrie 4.0" in 2011, while the Spanish government created the federal initiative "Industria Connectada 4.0" in 2014. Thus, it will be interesting to assess if starting the I4.0 introduction three years earlier leads to significant differences for the participating SMEs as would be expected.

The following table summarizes the hypotheses examined through the questionnaire:

| Hypothesis | Description |
|------------|---|
| H1 | Smaller SMEs need more assistance for evaluation of I4.0 technologies than bigger SMEs. |
| H2 | Smaller SMEs need more assistance to assess I4.0 introduction costs and benefits than bigger SMEs. |
| H3 | Smaller SMEs need more assistance to formulate an I4.0 strategy than bigger SMEs. |
| H4 | SMEs that collaborate with a big company feel better prepared for I4.0 introduction and have a higher technology acceptance rate than SMEs who do not collaborate with a big company. |
| H5 | Employees from SMEs with internal motivation to introduce I4.0 have a higher technology acceptance rate than employees from SMEs with external motivation to introduce I4.0. |
| H6 | SMEs with internal motivation to introduce I4.0 expect a higher increase in productivity than SMEs with external motivation to introduce I4.0. |
| H7 | There is a significant difference in the answers between Spanish and German SMEs. |

Table 2 - Industry 4.0 Questionnaire for SMEs Hypothesis

4.4 RESULTS

In the period from February to April 2020 overall 14 companies participated in the survey, with 11 completing the questionnaire. Seven of those companies were from Germany (63.6 %) and four (36.4 %) from Spain. The companies were also grouped by the size of their workforce: “less than 10 employees” (9.1 %), “10 to 49 employees” (18.2 %), “50 to 249 employees” (45.5 %) and “more than 250 employees” (27.3 %).

The Mann-Whitney U test was used to evaluate the results and find statistical significant differences in the answers of different groups (Mann & Whitney, 1947). The test was independently developed in 1947 by Mann-Whitney and Wilcoxon, thus it is also known as Wilcoxon-Mann-Whitney rank sum test. It is one of the most commonly used non-parametric tests, which means it does not depend on a normal distribution and provides reliable, statistically significant results when used with small sample sizes of 10-20 observations (Landers, 1981). The Mann-Whitney test verifies the null hypothesis (H_0) based on the comparison of each observation from the first group with each observation from the second group and identifies if the two independent groups are homogenous and have the same distribution (Nachar, 2008). If there is a statistically significant difference between the two populations, the determined p-value is small and the H_0 is rejected in favor of the alternative hypothesis. The commonly accepted thresholds are $p \leq 0.05$ for significant differences and $p \leq 0.01$ for highly significant differences and also used for this analysis (Fisher, 1992). All results stem from the two-sided test, which examines both ends of the distribution. Anova and similar techniques have not been used as normal distribution could not be guaranteed and the sample size is too small. Only the relevant questions to the particular hypothesis are shown.

The analysis of the questionnaire responses is done via the Python programming language¹³ (version 3.7). Within Python the following libraries have been used:

- The Pandas¹⁴ data analysis library was used to calculate general statistics (version 0.25.0)
- Matplotlib¹⁵ (version 3.1.1.) and the related Seaborn¹⁶ library (version 0.9.0) create the plots
- Scipy¹⁷, a library for scientific computing, provides an implementation of the Mann-Whitney U test (version 1.3.1)

H1: Smaller SMEs need more assistance for evaluation of I4.0 technologies than bigger SMEs.

Question Q16: Technology usage in companies.

Statistically significant differences in technology usage (Q16) comparing “Companies with up to 10 employees” with the other participants for all of the eight technologies: Sensor technology (p=0.009, highly significant), mobile end devices (p=0.004, highly significant), RFID (p= 0.027), real-time location systems (p=0.030), big data (p=0.018), cloud technologies (p=0.011), embedded IT systems (p=0.027) and M2M communication (p=0.022).

Splitting the participants into “Companies with up to 49 employees” and “others” identified significant differences for the following three out of eight technologies: Sensor technology (p=0.014), mobile end devices (p=0.006, highly significant) and big data(p=0.038).

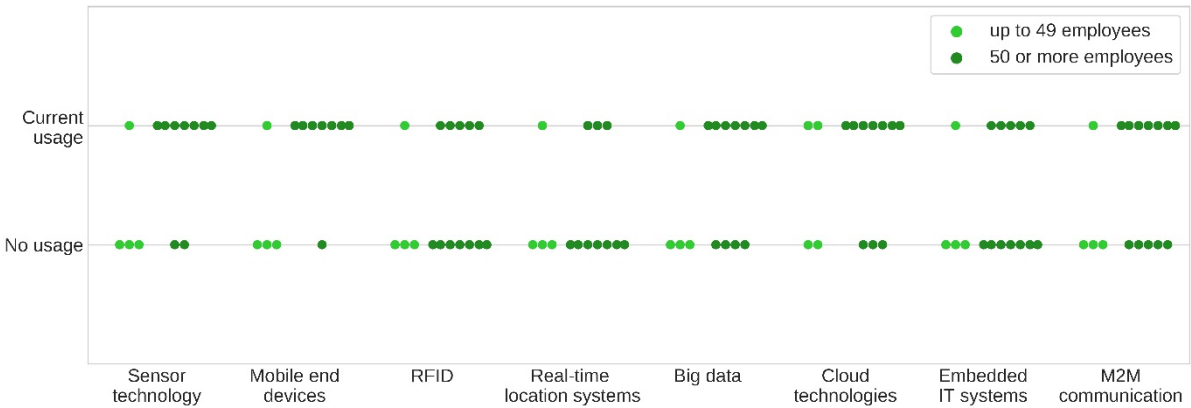


Figure 10 - Technology usage in companies with up to 49 vs more than 49 employees.

¹³ Python website: <https://www.python.org/> retrieved 18.11.2020
¹⁴ Pandas website: <https://pandas.pydata.org/> retrieved 18.11.2020
¹⁵ Matplotlib website: <https://matplotlib.org/> retrieved 18.11.2020
¹⁶ Seaborn website: <https://seaborn.pydata.org/> retrieved 18.11.2020
¹⁷ Scipy website: <https://scipy.org/> retrieved 18.11.2020

Question Q17: Past and future investments

The analysis of investments in the past 2 years found significant differences for “Research and development” ($p = 0.020$) and “Production / manufacturing” ($p = 0.023$).

The planned investments over the next 5 years also show significant differences in “Production / manufacturing” ($p = 0.035$). The analysis shows several significant differences and confirms H_1 .

H2: Smaller SMEs need more assistance to assess I4.0 introduction costs and benefits than bigger SMEs.

Analysis of Q12 “The benefits of I4.0 introduction are well known to our company and clearly evaluated” and Q13 “The costs of I4.0 introduction are well known to our company and clearly evaluated” yielded no significant differences between companies of different sizes. Thus H_2 is rejected.

H3: Smaller SMEs need more assistance to formulate an I4.0 strategy than bigger SMEs.

Question Q14: Industry 4.0 strategy implementation status

Significant results for all splits, i.e. “Companies with up to 10 employees” | “others” ($p = 0.030$), “Companies with up to 49 employees” | “others” ($p = 0.012$) and “Companies with up to 249 employees” | “others” ($p = 0.008$, highly significant).




Figure 11 - Q14 Industry 4.0 implementation status comparing companies with up to 49 employees and companies with more than 50 employees

Question Q15: Industry 4.0 indicators

Statistically significant differences for all splits, i.e. “Companies with up to 10 employees” | “others” ($p = 0.029$), “Companies with up to 49 employees” | “others” ($p = 0.010$) and “Companies with up to 249 employees” | “others” ($p = 0.008$, highly significant)




Figure 12 - Industry 4.0 indicators comparing companies with up to 49 employees and companies with more than 50 employees.

Question Q18: Systematic technology and innovation management

Split with “Companies with up to 10 employees” | “others” found significant differences in the technology and innovation management for IT ($p = 0.025$), production technologies ($p = 0.029$), product development ($p = 0.018$), services ($p = 0.031$) and in the amount of centralized innovation management ($p = 0.031$).

The split between “Companies with up to 49 employees” and “others” yields highly significant results ($p < 0.01$) for production technologies ($p = 0.003$) and product development ($p = 0.001$). Significant differences were found for services ($p = 0.044$) and the implementation of a centralized innovation management ($p=0.015$). The analysis shows significant or even highly significant differences and verifies H_3 .




Figure 13 - Systematic technology and innovation management in companies with up to 49 and more than 50 employees.

H4: SMEs that collaborate with a big company feel better prepared for I4.0 introduction and have a higher technology acceptance rate than SMEs who do not collaborate with a big company.

The samples are split based on their answer to question Q8 "Our company adopts the I4.0 strategy of a (bigger) partner".

Significant differences exist for Q7 "Our company is well prepared to introduce I4.0" ($p = 0.042$) and Q27 "Our employees face the new I4.0 challenges with confidence" ($p = 0.042$), thus H_4 is accepted.




Figure 14 - Comparing companies that agree/disagree to Q8 (mean + std.)

H5: Employees from SMEs with internal motivation to introduce I4.0 have a higher technology acceptance rate than employees from SMEs with external motivation to introduce I4.0.

The samples are split based on their answer to question Q10 "Our company feels the need to introduce I4.0 to continue collaboration with (bigger) partners". The split was chosen, as the other possible splits based on Q6 "Introducing I4.0 in our company is a good idea" and Q9 "Our company feels the need to introduce I4.0 to stay competitive" had uniformly agreeing answers.

Results are shown in an overview but are not statistically significant and therefore H_5 is rejected.

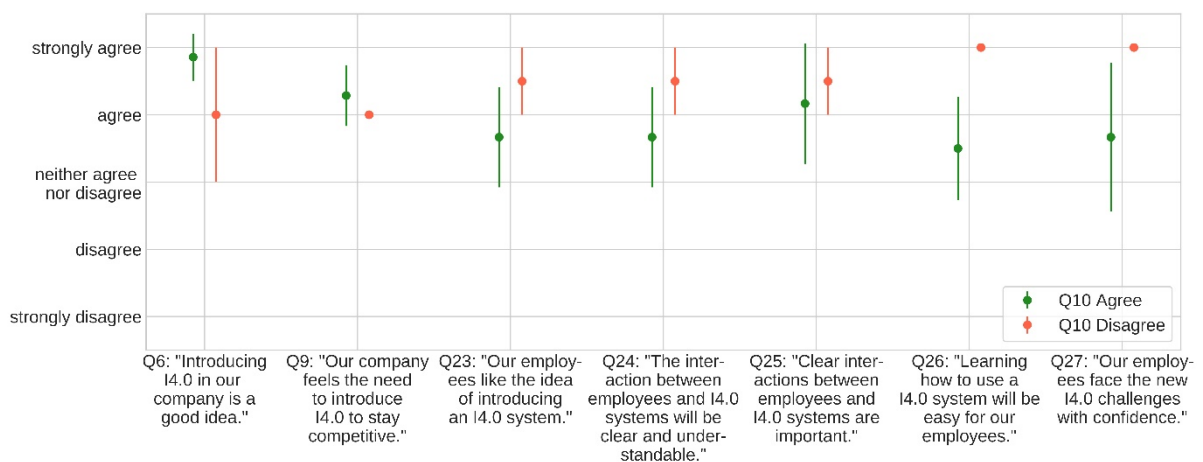


Figure 15 - H_5 overview with companies grouped based on their Q10 answers (mean + std.)

H6: SMEs with internal motivation to introduce I4.0 expect a higher increase in productivity than SMEs with external motivation to introduce I4.0.

The samples are split based on their answer to question Q10 "Our company feels the need to introduce I4.0 to continue collaboration with (bigger) partners.". The split was chosen, as the other possible splits based on Q6 "Introducing I4.0 in our company is a good idea" and Q9 "Our company feels the need to introduce I4.0 to stay competitive" had uniformly agreeing answers.

Results are shown in an overview but are not statistically significant and therefore H₆ is rejected.




Figure 16 - H₆ overview with companies grouped based on their Q10 answers (mean + std.)

H7: There is a significant difference in the answers between Spanish and German SMEs.

While 60 % of German SMEs planned to increase the employees from leadership working on I4.0 introduction, none of the Spanish SMEs plan additional workers (p = 0.007, highly significant). The situation is similar for an increase of employees in HR working on I4.0 introduction. 40 % of German companies plan to increase the current number of employees and none of the Spanish SMEs (p = 0.038).

A significant difference (p = 0.050) was also found between Spanish and German companies for Q9 "Our company feels the need to introduce I4.0 to stay competitive". Spanish companies tended to strongly agree (avg.: 4.75, std.: 0.43), while German companies agreed (avg.: 3.9, std.: 0.53).

The current status of I4.0 implementation (Q14, p = 0.016) and the indicators used to track the progress (Q15, p = 0.013) also differed significantly between the two countries. While half the German companies stated, that they have indicators that give them some orientation and just 10 % find their indicators to be already appropriate, 25 % of the Spanish companies have indicators that give them an orientation and 25 % think their indicators are already appropriate.

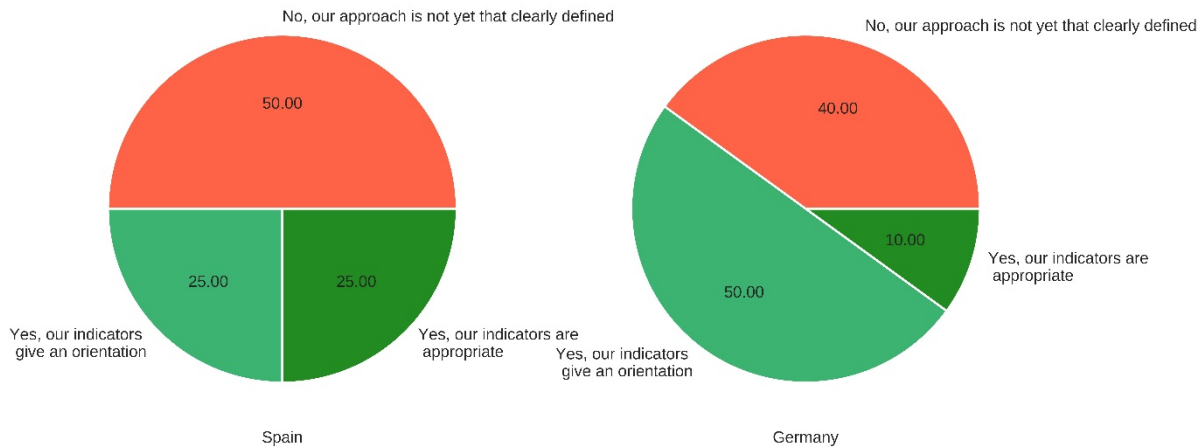


Figure 17 - Q15 I4.0 indicators with companies grouped by country

Surveying the existing technologies (Q16) in companies of both countries showed significantly more usage of mobile end devices ($p = 0.025$) in Germany while companies in Spain utilized more real-time location systems ($p = 0.030$). Unfortunately, there was also a highly significant difference of Spanish companies that use none of the inquired technologies ($p = 0.002$).

The results for technology and innovation management (Q18) also highlight differences between the two countries. The German companies focus on innovation management for production technologies ($p = 0.038$) and product development ($p = 0.012$) while the Spanish companies possess significantly more innovation management for their services ($p = 0.008$, highly significant) or use a centralized approach ($p = 0.002$, highly significant). The Spanish participants also declared significantly more companies without any technology or innovation management ($p = 0.040$). As the analysis found several statistical significant differences H_7 is accepted.

4.5 DISCUSSION

The results of the last section are summarized in Table 3. H_{1-3} regard additional assistance required by smaller companies to formulate an I4.0 strategy (H_1), assess costs and benefits (H_2) and evaluate the related technologies (H_3). H_1 and H_3 could be validated with significantly less technologies used, less systematic technology management, less investments made and also earlier stages of I4.0 introduction for smaller companies. Those findings may confirm claims by Christian Schröder that SMEs often lack resources to evaluate new technologies, which makes the development of an I4.0 strategy harder, and suggest that SMEs could not catch up regarding integration of their production plants in higher level IT systems since the survey by Lichtblau et al., which is a precondition for many I4.0 use cases (Lichtblau et al., 2015; Schröder, 2016). Apart from those differences all companies declared that they are able to evaluate the benefits but have problems to assess the associated costs of I4.0 introduction, thus H_2 is rejected.

| Hypothesis | Description | Result |
|------------|---|-------------------------------------|
| H1 | Smaller SMEs need more assistance for evaluation of I4.0 technologies than bigger SMEs. | <input checked="" type="checkbox"/> |
| H2 | Smaller SMEs need more assistance to assess I4.0 introduction costs and benefits than bigger SMEs. | <input type="checkbox"/> |
| H3 | Smaller SMEs need more assistance to formulate an I4.0 strategy than bigger SMEs. | <input checked="" type="checkbox"/> |
| H4 | SMEs that collaborate with a big company feel better prepared for I4.0 introduction and have a higher technology acceptance rate than SMEs who do not collaborate with a big company. | <input checked="" type="checkbox"/> |
| H5 | Employees from SMEs with internal motivation to introduce I4.0 have a higher technology acceptance rate than employees from SMEs with external motivation to introduce I4.0. | <input type="checkbox"/> |
| H6 | SMEs with internal motivation to introduce I4.0 expect a higher increase in productivity than SMEs with external motivation to introduce I4.0. | <input type="checkbox"/> |
| H7 | There is a significant difference in the answers between Spanish and German SMEs. | <input checked="" type="checkbox"/> |

Table 3 - Industry 4.0 Questionnaire for SMEs Results

The collaboration with a bigger partner on the I4.0 introduction led to a significantly more positive attitude towards I4.0, which confirmed H₄. Even though “many of the I4.0 methods are developed mainly in larger enterprises”, there is potential for the SMEs to profit from the groundwork done by the bigger partner, especially when the SMEs own resources are rather scarce (Rauch et al., 2018). Five years after Lutz Sommer stated that “most SMEs are not prepared to implement I4.0 concepts” the collaboration with a bigger partner leads to SMEs who feel well prepared for the I4.0 introduction (Sommer, 2015).

H₅₋₆ examine the influence of internal and external motivation to introduce I4.0 towards the technology acceptance rate and expected increases of productivity but could not be statistically verified, thus both hypothesis are rejected.

The comparison of Spanish and German SMEs highlighted various statistically significant differences which leads to acceptance of H₇. Most noticeably are Spanish companies, that use none of the new technologies and also do not possess technology or innovation management. It is not possible to determine if it is a regional difference or if it is caused by the small sample size where Spanish companies are smaller on average. The small sample size is the main limitation of the questionnaire and stems from the specific requirements for the participants, but also the world-wide COVID-19 pandemic where SMEs had to shut down their production. It would be interesting to conduct the questionnaire again with more participants to get more insights in the differences between Spanish and German SMEs, even though the results are already statistically significant.

4.6 CONCLUSION

Conducting the Industry 4.0 Questionnaire for SMEs with Spanish and German participants provides an insight into the status of the I4.0 implementation in the companies and their attitude towards the new technologies. The results allow the following conclusions:

- Smaller SMEs need more assistance for evaluation of I4.0 technologies than bigger SMEs.
- Smaller SMEs need more assistance to formulate an I4.0 strategy than bigger SMEs.
- SMEs that collaborate with a big company feel better prepared for I4.0 introduction and have a higher technology acceptance rate than SMEs who do not collaborate with a big company.

Future research should work on concrete methods to assist SMEs in the development of an I4.0 strategy and the evaluation of the associated new technologies.

The majority of participants also stated that they need assistance to formulate an Industry 4.0 strategy and evaluate the available Industry 4.0 technologies. They are looking for standards or best practices, which unfortunately are not yet available, as stated by Matt and Rauch, will provide great value to SMEs (D. Matt & Rauch, 2020). Reference architectures, presented in the next section, support SMEs in the I4.0 introduction through a domain-specific template for the implementation of CPPS software systems.

5 REFERENCE ARCHITECTURES FOR INDUSTRY 4.0

A reference architecture (RA) supports complex software systems' design through the abstraction of software systems for a specific domain. Software systems for CPPS are often complex as they incorporate the requirements from mechanical and electrical engineering as well as computer science. Thus, the support from a RA is valuable for companies of any size, but especially for small- and medium-sized companies, which often do not have the capacities or budget to build machine learning systems for CPPS from scratch.

Our workgroup created a catalog of requirements and evaluated proposed reference architectures from the field of automation and cognition on three real-world use cases to highlight the existing gap. The results have been presented at the 24th IEEE International Conference on Emerging Technologies and Factory Automation in Zaragoza (Bunte, Fischbach, et al., 2019).

The remainder of this chapter is based on and extending that work. Thus, at first, RAs are introduced more in-depth. Then different use cases and the resulting consolidated requirements are presented. Afterwards, the proposed reference architectures are portrayed and finally evaluated on the derived requirements.

5.1 REFERENCE ARCHITECTURE

A RA supports the design of complex software systems, e.g., implementing artificial intelligence algorithms for a CPPS in the context of the work at hand.

A key concept for RAs is the abstraction of concrete systems to extract their essence and describe a class of systems (Angelov et al., 2012; Nakagawa et al., 2012). Thus, the RA divides the required systems' functionality to solve a domain-specific class of problems into separate components and indicates the information flow between those components (Bass et al., 1998; IEEE Computer Society, 2000; Vogel et al., 2011).

Hu et al. provide a straightforward architecture for the design of CPPS, which highlights the abstraction and is shown in Figure 18 (Hu et al., 2012). They capture the essence of a CPPS. The physical machine senses the environment and derives a new control decision from the information via a network connection to the service framework and the control system. This decision is executed through the physical machines' actuators and translated back into the target environment.

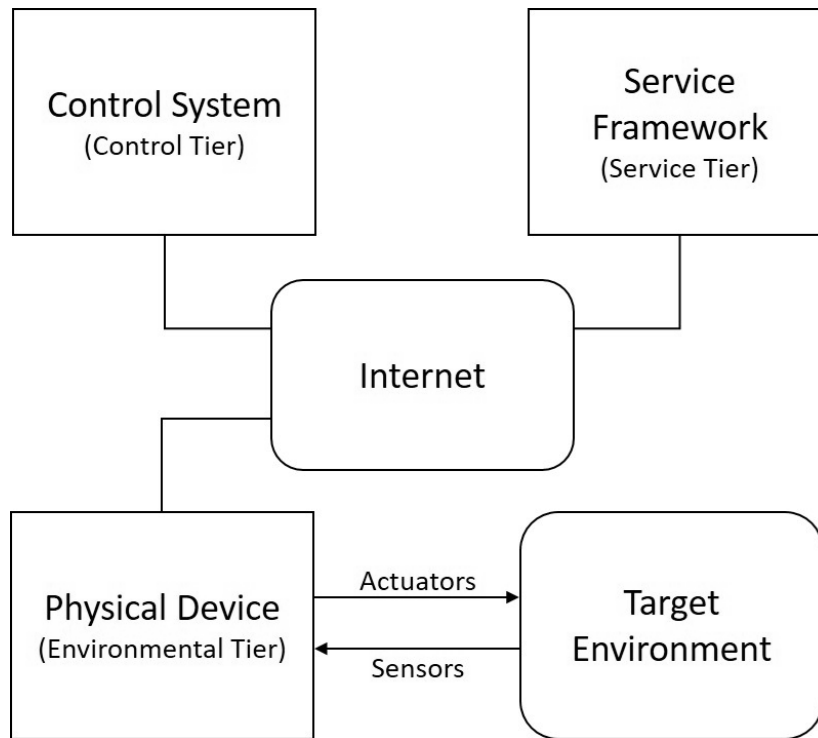


Figure 18 - CPPS Architecture

While this architecture captures the idea of a CPPS, it may be too abstract to support the implementation of a CPPS software system. Finding the appropriate level of abstraction is one of the main challenges for the creation of a new reference architecture. The RA has to be concrete enough to provide sufficiently specific information and guidelines to implement a system, but not too big as then the “essence is hidden instead of highlighted” (Cloutier et al., 2010).

Thus, a good RA makes the implicit domain knowledge explicit, e.g., through the utilization of design patterns, and guides future system development without losing too many details (Campetelli et al., 2014). The resulting RA uses domain knowledge to fulfill the specific requirements but is still general enough for a flexible application in different use cases within this domain (M. Maier, 2013).

5.1.1 Reference architecture benefits

Van Brussel et al. developed the PROSA reference architecture for traditional manufacturing systems (Van Brussel et al., 1998). They found reference architectures to be useful for the manufacturing domain as the decoupling of components “enables an intensive reuse of sub-systems in a wide range of manufacturing systems”, which should lead to increased reliability of the system, improved performance and reduced installation and maintenance costs.

Later Cloutier et al. compiled a general overview of RA benefits (Cloutier et al., 2010). RAs provide:

- a common lexicon and taxonomy
- a common (architectural) vision
- modularization and the complementary context

Thus, the researchers concluded that RAs are useful for:

- managing synergy
- providing guidance, e.g., architecture principles and best practices
- providing an architecture baseline and an architecture blueprint
- capturing and sharing (architectural) patterns on the basis of proven concepts within a domain

Cloutier et al. asked the participants from the “Systems Architecture Forum” to identify the key points for the application of RAs in industry. They found the following motives:

- Reuse and commonality of systems or parts of systems which lead to shorter development cycles and reduced costs as implementation does not start from scratch
- Risk reduction and higher quality through the use of proven patterns and architectural elements
- Increased interoperability and compatibility between systems but also system vendors and suppliers through standardization
- RA as a knowledge repository facilitates communication and knowledge transfer
- Controlling complexity through composition of well-known and practiced patterns

5.1.2 Developing a reference architecture

Galster and Avgeriou suggest that a new reference architecture should be “empirically-grounded”, i.e. build on an “empirical foundation” with “empirical validity” (Galster & Avgeriou, 2011). The empirical foundation consists of three aspects:

- The RA addresses real stakeholders’ interests
- The RA is based on concepts that have been proven in practice
- The RA uses building blocks from the problem domain

Empirical validity means that a RA must prove its applicability and validity for a given domain.

A resulting empirically-grounded RA shall be applicable not only in one specific organization, but instead for a broad range of use cases.

Galster and Avgeriou further note that the development of a RA can be either “practice-driven” or “research-driven”. A practice-driven RA can be defined when the existing knowledge on a domain allows to propose best practices. The research-driven RA in contrast is “inspired by existing research effort and provides a futuristic view of a class of systems”.

Angelov et al. extend these architecture types based on when, why and where they are created (Angelov et al., 2012). They named the practice-driven RA a “classical” RA as it is developed after enough experience in a specific domain has been accumulated. The research-driven RA is considered a “preliminary” RA as it is created before any system fully implements the architecture. They further partitioned architectures based on the goal of the architecture in “standardization” and “facilitation” RAs, where a standardization RA focuses on the interoperability of systems created with the same RA and a facilitation RA aims to support the design of systems. The final distinction regards the audience, whether the RA is created for a single or multiple organizations.

The goal for the CAAI architecture is to be a research-driven facilitation RA that introduces artificial intelligence in the CPPS with a focus on easy implementation for multiple organizations.

5.1.3 Validating a reference architecture

The value of a RA can be assessed based on its correctness and utility within a project and the support for efficient implementation and adaptation (Galster & Avgeriou, 2011). However, the validation of a research-driven RA that introduces disruptive technologies or innovative applications can be difficult and in this case an iterative development approach is recommended (Cloutier et al., 2010; Kruchten, 1995). Actually implementing a prototype of the architecture allows to test, analyze and refine the architecture in subsequent iterations. Additionally the iterative approach aids the development of a research-driven architecture as it helps to further refine the architectural requirements. In this regard Roberts and Johnson created the “rule of three” which claims that it is necessary to develop “at least three independent, observable applications” to generalize successfully and find hidden abstractions (Roberts & Johnson, 1997).

Another approach to the validation of a RA is the usage and verification through other developers, as a RA becomes more valuable the more people use it. Brad Appleton points out that a proposed solution needs to be investigated and scrutinized by others before it can become a reference (Appleton, 2000). Similarly Cloutier et al. report that this scrutiny gives potential users “confidence in the completeness and appropriateness” of a solution (Cloutier et al., 2010).

Thus, the CAAI is developed in several iterations and all results are published as open-source software¹⁸.

5.1.4 Need for a cognitive reference architecture for CPPS

Campetelli et al. investigated the use of reference architectures for CPS (Campetelli et al., 2014). They found a wide variety of application domains, e.g., automotive, chemical processes, medical instruments or manufacturing, and propose to develop reference architectures for each of those heterogeneous domains.

A software system for CPPS that introduces AI into manufacturing has to process vast amounts of data. Thus, it was an obvious choice to evaluate big data reference architectures as the foundation for further developments.

The Lambda architecture and its successor the Kappa architecture are designed to handle massive quantities of data (Kreps, 2014b, 2014a; Marz & Warren, 2015). Data processing is divided into batch and stream processing in the Lambda architecture. The advantage of this approach is to use the best implementation for each, the processing of big data sets in scheduled intervals and the continuous processing of single data points in real-time. In practice the maintenance of two codebases turned out to be a great effort. Thus, the Kappa architecture unifies batch and stream processing and points out that a batch is also a stream, just with a defined start and end point. At the same time Pääkkönen & Pakkala compared the existing big data implementations from Facebook, Twitter, LinkedIn and Netflix to derive a more detailed reference architecture for big data applications (Pääkkönen & Pakkala, 2015).

While the event-driven approach of all those architectures is well suited for CPPS use cases, these architectures, of course, do not take into account the special requirements of intelligent manufacturing systems. The CPPS, as defined in 1.3.1, is a "system of collaborating computational entities which are in intensive connection with the surrounding physical world and its on-going processes, providing and using, at the same time, data-accessing and data-processing services" (Monostori et al., 2016). Sensing

¹⁸ <https://github.com/janstrohschein/KOARCH>

the environment, processing the information and initiating actions that influence the environment are cognitive capabilities that need to be considered.

Neisser defines cognition in his works on cognitive psychology as “all processes by which the sensory input is transformed, reduced, elaborated, stored, recovered and used” (Neisser, 2014). This definition can be adapted for the I4.0 context: Cognition for a CPPS refers to all processes which transform, reduce, elaborate, store, recover and use the input data to solve I4.0 use cases, i.e., condition monitoring, anomaly detection, optimization and predictive maintenance.

Thus, a cognitive CPPS learns from the sensor information and independently adapts to changing environments. Cognitive capabilities enable skills such as learning, planning and decision making, e.g. production planning based on incoming orders, optimizing production and detecting anomalies before a component fails. Ultimately, a cognitive reference architecture should be developed.

The data processing capabilities of big data architectures will be the foundation for a cognitive reference architecture and extended for the special requirements of I4.0 and CPPS. The next sections present the special requirements and evaluate proposed reference architectures for their suitability.

5.2 USE CASES AND REQUIREMENTS

A RA for I4.0 needs to assist companies in the implementation of a CPPS. Thus, real-world requirements are derived from real-world use cases to collect the demands towards a RA in the context of I4.0. A consolidation of the requirements from different use cases highlights the essence of such systems and allows the abstraction into a reference architecture.

5.2.1 Requirements A: Diagnosis for a modular CPPS

The VPS in the SmartFactory-OWL¹⁹ is a demonstration plant of a modular CPPS. The VPS, shown in Figure 19 below, processes corn in different modules and finally produces popcorn.

It consists of the modules delivery, storage, quality control, dosing and production. Modules have compatible interfaces, which allow different hardware configurations to enable adaptive production. It is also possible to extend the system with new modules, which might be unknown during the engineering phase of the initial modules. Therefore a self-diagnosis system is needed to detect anomalies independently of the current configuration, determine the root-cause and stop the faulty modules to prevent additional damage. AI algorithms enable a model-based diagnosis as seen in the publication by Bunte, Stein and Niggemann (Bunte, Stein, et al., 2019).

¹⁹ SmartFactory-OWL Website: <https://smartfactory-owl.de/?lang=en>



Figure 19 - Versatile Production System, picture by Andreas Bunte

This use case is an example of value based services, as described in the Condition Monitoring scenarios through the Plattform Industrie 4.0 (Bauer et al., 2017). However, the self-reaction of the CPPS goes beyond that use case. Due to the versatile combination possibilities, no diagnosis system is implemented in the VPS, since approaches known today cannot handle these versatile systems efficiently. Therefore, data driven approaches can be an efficient method for anomaly detection, as shown by Maier et al (A. Maier et al., 2015). These methods provide a huge potential, but their implementation is difficult, because there are no common interfaces and thus it has to be adapted for every specific application with a certain effort. The architecture must address the following requirements to support the implementation:

- (R-A.1) Data acquisition via Open Platform Communications Unified Architecture (OPC UA) server and Modbus.
- (R-A.2) Store data to collect all lifecycle information, such as process data and models of the CPPS.
- (R-A.3) Perform preprocessing to handle missing values, normalize data or adapt the format.
- (R-A.4) Learn diagnosis models that enable diagnosis with few learning data sets for online and offline learning.
- (R-A.5) Provide a diagnosis algorithm that performs diagnosis in less than 100 ms.
- (R-A.6) Decision making whether or not a response is required and choosing an action.
- (R-A.7) Access the controller to perform the chosen action and thus prevent further damage.

5.2.2 Requirements B: Energy Efficiency Optimization in Bakeries

Optimizing the usage of energy is an important task especially in industries with a high energy consumption. Consider bakeries that consists of several chain stores of different types, i.e., sale only, production only and stores combining both, production and sale. Production devices, especially ovens, are the major energy consumption devices. Ovens contain several trays, which can be controlled individually. Planning an optimal baking procedure with different products at different temperatures is a crucial task w.r.t. energy efficiency, as idle times for the ovens and unnecessary cool-down times

are to avoid. Companies energy costs are not only calculated by the consumption, also the maximum used power has to be paid. Thus for companies it is beneficial to divide the energy consumption over the day, instead of creating large peaks by, e.g., turning on all devices at start of business. Thus a smart start up schedule provides a significant contribution in cost optimization. Furthermore stores typically also have a large divergence regarding the numbers of arriving customers at different times during day, and thus necessary stock of products.

Additionally to the monitoring of energy consumption, the predicted amount of products to sell and the available stock of products should be regarded. A smart integration of the inventory and sale system enables the computation of accurate decentralized models of product sales for a given time of each day. Enabling information exchange between different models can lead to a higher level of adaptability. The architecture must address the following requirements:

- (R-B.1) Highly distributed and heterogeneous systems.
- (R-B.2) Store data, such that suitable historical data can be used, e.g., for optimization.
- (R-B.3) Preprocess data, e.g., treat missing data accordingly.
- (R-B.4) Learn and update models of energy consumption.
- (R-B.5) Provide a simulation to enable energy efficiency optimization and peak prevention.
- (R-B.6) Provide guidance to employees by implementing an appropriate human machine interface (HMI).
-

5.2.3 Requirements C: Process Control of Concrete Spreading Machines

A concrete spreading machine produces pre-cast concrete components as shown in Figure 20 based on the work by Heienbrok (Heienbrok, 1997).




Figure 20 - Concrete spreading machine with casings

The mold consists of a steel pallet, casing and additional reinforcements. The machine pours concrete on the steel pallet. Casing and reinforcements mounted on the steel pallet determine the shape and properties of a component. Controlling parameters for the process have to be set manually, which is quite difficult and needs a lot of experience since some parameters are hard to assess, e.g., the consistency of concrete from an earlier cast mixed with fresh concrete. Additionally changing parameters will not lead to an immediate reaction as the concrete moves slowly. The goal of this use case is to learn a model that is able to control the process in order to optimize two conflicting goals: minimizing the production time while maximizing the quality of the resulting product. In order to implement a system that learns from process data to control and optimize the production output the following requirements towards an architecture emerge:

- (R-C.1) Record sensor data and perform preprocessing to extract standardized process information.
- (R-C.2) Store relevant information of the production process.
- (R-C.3) Learn models from historical as well as current sensor data to control the process and consider conflicting goals.
- (R-C.4) Decision making in real-time or near-real-time to be useful in a real-world production process.
- (R-C.5) Verify results in a simulation before allowing the algorithm to control an actual concrete spreading machine.
- (R-C.6) Communicate with the machine in a standardized format for compatibility with different models/vendors.

5.2.4 Requirements D: General Requirements

The aim of the cognitive reference architecture is to cover versatile use cases within one architecture. Therefore implementations can be easily adapted to different use cases which reduces engineering costs. In a specific use case, the adaption might be the selection of a proper algorithm, based on the user specification. The architecture must enable the users to implement decision making and to learn which decisions are promising for which application. To achieve this, there are some additional requirements that cannot be derived from the single use cases, so there are some overall requirements:

- (R-D.1) Receive declarative goals from the user.
- (R-D.2) Specified interfaces are well defined.
- (R-D.3) Strategies to select a suitable algorithm.
- (R-D.4) The system learns from experiences.
- (R-D.5) Thorough knowledge representation.

The user should be able to define goals easily. If the system accepts declarative goals (R-D.1) it is possible to define a goal as plain as "optimize energy consumption". Additionally the user can specify additional constraints, such as limitation of the response time. Furthermore the reference architecture should possess well defined interfaces (R-D.2) with a thorough description of which information should be transferred to ensure modular expandability. Re-usability and customization benefit from such a modular structure and this will also enable concepts such as software as a service. Additionally exchange and purchase of services between vendors are possible. The reference architecture needs a

strategy to select an algorithm (R-D.3) that is able to produce a feasible result under the current conditions, such as volume of data, the required response time or the type of problem to solve. A cognitive architecture reflects upon the decision making and is able to learn from past experiences (R-D.4) to increase processing efficiency and outcome over time. Essentially what was the right tool for a given job. To aid this decision making and to model/store additionally learned knowledge it is required to implement a suitable knowledge representation (R-D.5) that represents information about machines and processes, newly learned rules and domain knowledge from experts.

5.2.5 Requirement Consolidation

The requirements from the three real-world use cases (A-C) and the general requirements (D) can be aggregated into a consolidated list of requirements towards a cognitive reference architecture for I4.0. The derived requirements in *Table 4* display the intersection of requirements from different use cases along with a requirement description.

| Requirement | Source | Description |
|-------------|---------------------|--|
| R.1 | R-D.1 | Receive declarative goals from the user |
| R.2 | R-D.2 | Specified interfaces are well defined |
| R.3 | R-D.3 | Strategies to select a suitable algorithm |
| R.4 | R-D.4 | The system learns from experiences |
| R.5 | R-D.5 | Knowledge representation including expert knowledge |
| R.6 | R-A.1, R-B.1, R-C.1 | Acquire data from heterogeneous and distributed systems via versatile protocols |
| R.7 | R-A.2, R-B.2, R-C.2 | Store and manage acquired process data and models |
| R.8 | R-A.3, R-B.3, R-C.1 | Perform preprocessing |
| R.9 | R-A.4, R-B.4, R-C3 | Learn a model from data, might be time and resource limited |
| R.10 | R-A.5, R-B.5, R-C5 | Perform a model analysis or simulation to get valuable information, might have a limited response time |
| R.11 | R-B.6 | Interaction with the user, i.e., human-machine-interface |
| R.12 | R-A.6, R-C4 | Decision making to derive actions, e.g., decide to send new control parameters to the CPPS controller. |
| R.13 | R-A.7, R-C.6 | Apply the action on the controller |

Table 4 – Consolidated requirements for CPPS reference architectures in I4.0

Those requirements are used in the following sections to evaluate proposed reference architectures for the application in I4.0 use cases. The evaluation also shows the potential shortcomings of existing reference architectures.

5.3 OVERVIEW OF REFERENCE ARCHITECTURES FOR INDUSTRY 4.0

A cognitive RA for the application in I4.0 has to fulfill the consolidated requirements for a successful introduction in several and distinct use cases within a company. There are two promising architecture classes to implement cognitive CPPS, namely architectures from the field of automation and cognitive architectures from cognitive sciences. The RAs from automation present a generic structure for a class of architectures that support automation systems design. The selected candidate architectures from

automation are well-known, namely the RAMI4.0, IIRA, and the 5C architecture. The word cognitive is used on a higher level in automation than in cognitive sciences as fulfilling automation goals may not require fully understanding human cognition. A cognitive architecture in automation adapts the process independently to new situations, such as creating an adjusted production plan, if a module breaks. Compared to that, traditional cognitive architectures from cognitive sciences have the purpose of understanding and reproducing human cognition. Lehman et al. defined a cognitive architecture as “a theory of the fixed mechanisms and structures that underlie human cognition” (Lehman et al., 2006). The established architectures Soar and ACT-R are introduced as candidate architectures from this field.

5.3.1 Reference Architecture Model Industry 4.0

In 2013, the three associations Bitkom, VDMA and ZVEI founded the Industry 4.0 platform to bring together the different interests and requirements from the fields of mechanical engineering, electrical engineering and IT towards a joint model for I4.0. The resulting Reference Architecture Model Industry 4.0 (RAMI4.0) was published by the associations in 2015 and also submitted for standardization in 2016 as DIN SPEC 91345 (Bitkom et al., 2015; DIN, 2016; VDI, 2015).

Overview

The RA model, as shown in Figure 21 below, has three different dimensions to represent the essential aspects of I4.0, i.e., Layers, Life Cycle & Value Stream and Hierarchy Levels (VDI, 2015).




Figure 21 - RAMI4.0 overview

Layers

The six different layers are based on the Smart Grid Architecture Model (Smart Grid Coordination Group, 2012), which was introduced in 2012 by the Smart Grid Coordination Group and extended for I4.0 deployment. They represent the digital image of a machine, for example, and were described in 2016 as follows (Schleipen, 2016):

- **Asset:** Representation of reality, e.g. technical object
- **Integration:** Provision of the asset's information
- **Communication:** Standardization of communication, using a uniform data format
- **Information:** Runtime Environment for Event (Pre)Processing
- **Functional:** Runtime and modeling environment for services that support business processes
- **Business:** Business models and the resulting business process. Ensures the integrity of the functions in the value chain.

The layers are loosely coupled and events and data are meant to be exchanged within a layer or between adjacent layers (VDI, 2015).

Life Cycle & Value Stream

The left horizontal axis represents the life cycle of plants and products, based on IEC 62890 regarding life cycle management of industrial systems (DIN, 2017). A distinction is also made between type and instance. A "type" becomes an "instance" when the development and prototype production is completed and the actual product is manufactured in the production department (Hankel, 2015).

The VDI status report presents an example, shown in Figure 22 below (VDI, 2015). As the customer buys a manufactured product, it initially becomes a type again, since installation in a new plant is planned. It only becomes an instance again when it is actually installed. This change between type and instance can thus be repeated several times.

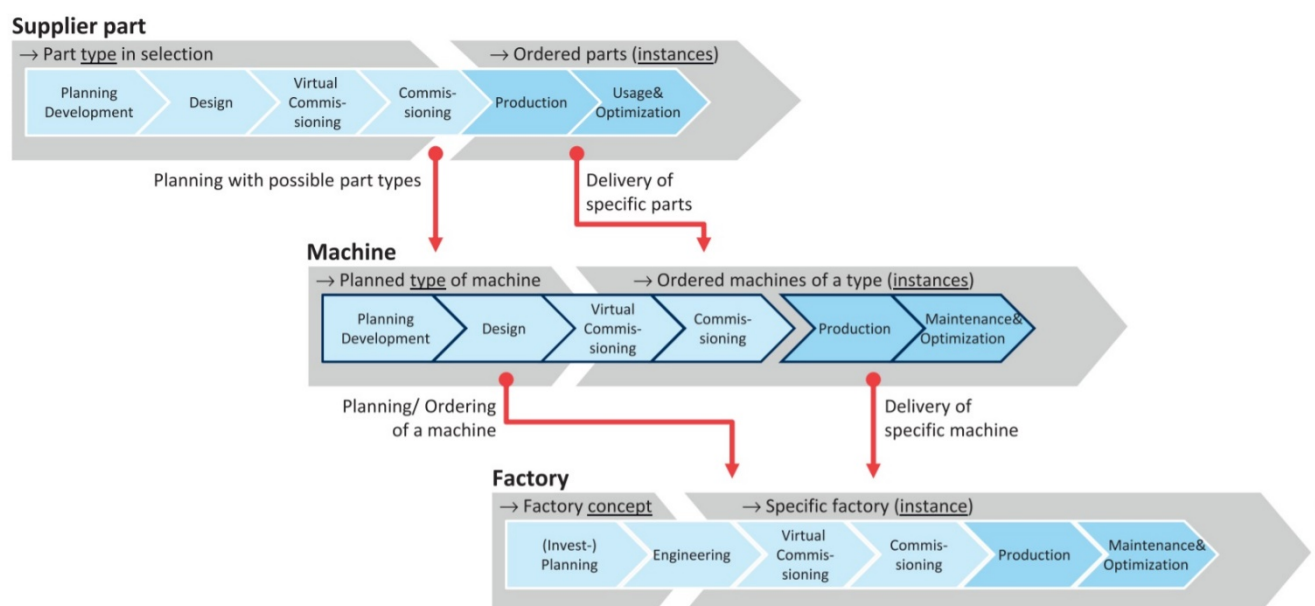


Figure 22 - Types and instances in a life cycle

Hierarchy Levels

The third axis of RAMI 4.0 describes the functional classification and responsibilities of a physical asset within I4.0. This is not about implementation, it is only about classification. For the classification within

a factory, the reference architecture model for this axis is based on the standards IEC 62264 and IEC 61512 (DIN, 2000, 2014) and extends the automation pyramid. For a uniform consideration across as many sectors as possible from process industry to factory automation, the terms "Enterprise", "Work Unit", "Station" and "Control Device" were adopted from the options listed there, while the "Field Device" was added. It represents the functional level of an intelligent field device, for example an intelligent sensor. Since the product to be manufactured is also important for the considerations, it was also included as "Product". I4.0 also describes the factory network, the cooperation with external engineering offices, suppliers and customers. This aspect is considered in the "Connected World", which has been inserted at the top of the hierarchy levels (Hübner, 2015).

The Industry 4.0 component

The I4.0 RA also derives a first definition for I4.0 components. An I4.0 component can represent a production system, a single machine or an assembly within a machine. An I4.0 component can consist of 1:n objects, but also of 1:n I4.0 components. In order for an object to be called an I4.0 component, it must have communication capabilities and a virtual representation. Existing objects that do not yet have this communication capability can be "extended" to become I4.0 components by having a higher-level IT system provide this function. The virtual image of an I4.0 component can also be provided either by the component itself or by a higher-level system, provided that I4.0-compliant communication with the outside world is possible (VDI, 2015). Throughout the complete life cycle of an I4.0 component, all relevant data should be stored in an electronic container and made available to the authorized parties. This container is called an "administration shell" and contains the virtual image of an I4.0 component. Figure 23 shows an I4.0 component that consists of several objects and has a management shell including manifest and resource manager (VDI, 2015). The resource manager enables access to data and functions of the management shell while the following metadata can be stored in the manifest according to VDI:

- Characteristic features of the real component
- Information on relations of the characteristics among themselves
- Production and production process relevant relationships between I4.0 components
- Formal description of relevant functions of the machine and its processes

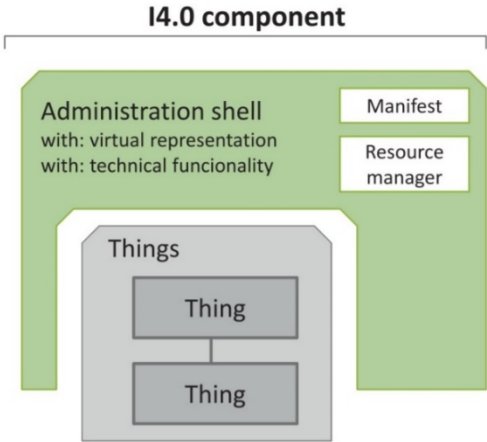


Figure 23 - Industry 4.0 component

In 2016 a document regarding the administration shell was published by ZVEI (ZVEI, 2016). It explains how the different aspects of an I4.0 component can be represented by sub-models and standardized individually. The most recent version contains an example with a drilling machine, which could have the sub-models "Drilling" and "Energy Efficiency" and describe each with the appropriate features in the administration shell (BMW, 2018).

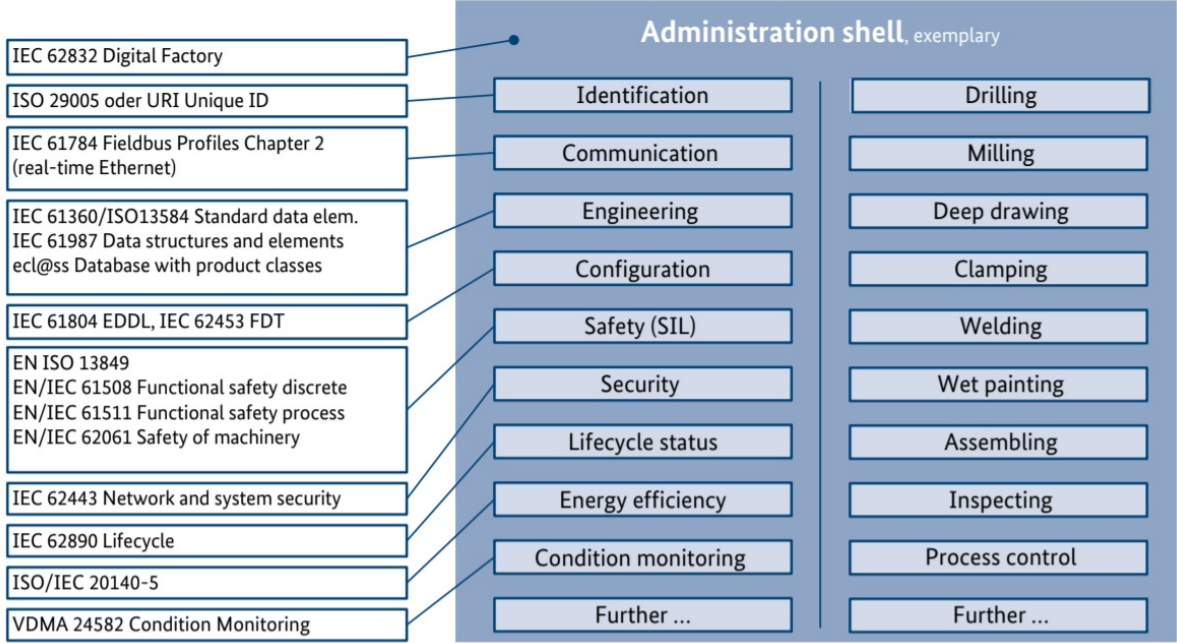


Figure 24 - Administration shell composed of relevant sub-models

RAMI 4.0 provides the basis for further steps. According to Martin Hankel and ZVEI, these are the unique identification of individual components in networked production, uniform semantics/syntax for manufacturer-independent data exchange between machines, a definition of the Quality of Services of the Industry 4.0 components and a standardization of the Industry 4.0 communication (Hankel, 2015). For some of these topics there are already proposed solutions, such as OPC UA, a protocol for communication between machines, but there is no reference solution that has been accepted as a standard by all industry partners.

Benefits and goals

The model unites the different user perspectives and creates a common understanding of industry 4.0 technologies. On the basis of RAMI 4.0, the requirements of the user industries - from production automation to mechanical engineering to process engineering - can be discussed in the corresponding committees of the associations and standardization bodies. The model thus creates a common understanding for standards, norms and practical case studies.

RAMI 4.0 serves as a map for industry 4.0 solutions: The model provides an orientation on which the requirements of the user industries are plotted together with national and international standards in order to define and further develop Industry 4.0. In this way, overlaps and gaps in standardization become visible and can be closed. (Hankel, 2015)

The goals for RAMI 4.0 are according to the document for the implementation strategy regarding Industry 4.0 (Bitkom et al., 2015):

- Clear and simple architectural model for reference
- Location of existing norms and standards
- Identification and closing of gaps in norms and standards
- Identification of overlaps and definition of preferred solutions
- Minimization of the number of norms and standards used
- Identification of subsets of a norm or standard for the rapid implementation of partial contents for industry 4.0 ("I4.0-Ready")
- Location of use case contents
- Identification of relationships
- Parent rule definition

RAMI4.0 can be used to classify machines and components of a CPPS and provides a standardized approach for the design of I4.0 machines. Current assets can be upgraded to I4.0 components as the additional functionality may be provided by an external system. However, in its current form, RAMI 4.0 does not support the concrete software design of an I4.0 CPPS.

5.3.2 Industrial Internet Reference Architecture

The IIC first introduced its Industrial Internet Reference Architecture (IIRA) in 2015 and presented the most recent version 1.9 in 2019 (S.-W. Lin et al., 2019). The architectural template is based on ISO/IEC/IEEE 42010 standard for architecture description and provides a methodology for IIoT System architects to design their systems based on a common framework and concepts (ISO, 2011). Figure 25 depicts the three architecture dimensions, i.e., Viewpoints, Lifecycle Process and Industrial Sectors.

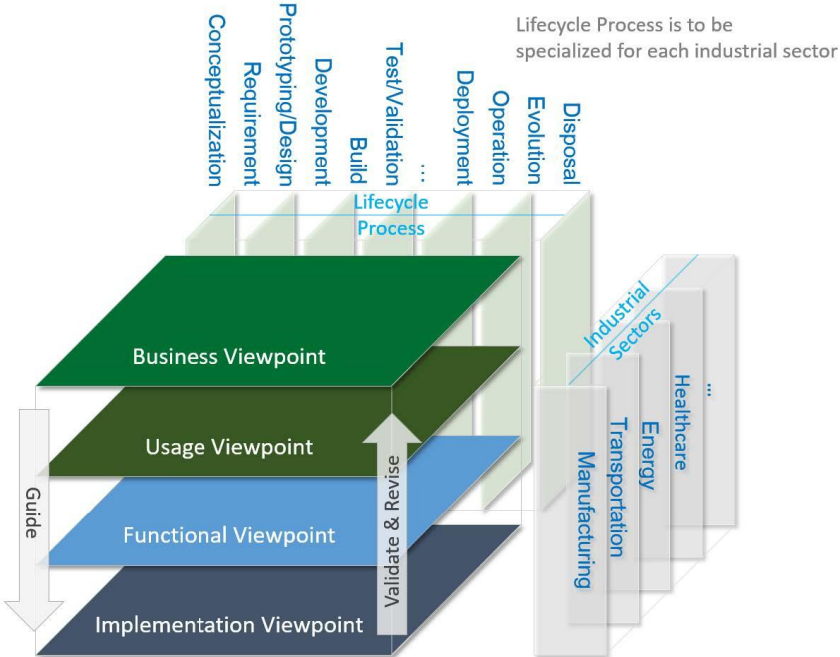


Figure 25 - Industrial Internet Reference Architecture Overview

Viewpoints

The IIC derived four architecture viewpoints, i.e., business, usage, functional and implementation, through analysis of various IIOT use cases:

- The business viewpoint identifies all relevant stakeholders, their business vision and objectives for the IIoT system. The stated objectives are mapped to fundamental system capabilities.
- The usage viewpoint represents the system usage through sequences of activities carried out by humans or computer systems to achieve the fundamental system capabilities.
- The functional viewpoint addresses the functional components that support the usage and activities of the overall system. The relationships and interactions between those components are defined with the respective interfaces.
- The implementation viewpoint concerns the required technologies to implement the functional components, their communication and life cycles.

The viewpoints are ordered in a top-to-bottom logic, where each viewpoint defines the requirements for the viewpoints below. Additionally, the viewpoints below validate the viewpoints in the layers above and indicate if the implementation is not feasible.

Lifecycle Process

The IIRA is not a description of a system lifecycle process, as the process varies from one industrial sector to another. It is a tool for system conceptualization that highlights important system concerns that may affect the lifecycle process. It guides system lifecycle processes from IIoT system conception to design and implementation through its viewpoints. The IIC plans to release additional documents regarding IIoT system lifecycle processes.

Industrial Sectors

The IIC wants to ensure a wide applicability of the IIRA for IIoT applications across all industrial sectors, thus the reference architecture is generic on purpose and the framework provides high level architecture descriptions and a high degree of abstraction. The related stakeholders in each viewpoints must regard the affected industrial sectors, to gather and describe the specific concerns resulting in the unique system requirements, thus transforming the abstract architectural concepts and models into a detailed architecture.

The IIRA summarizes common concerns from different perspectives of the stakeholders gathered from lots of use cases and projects, and therefore represents a high level of abstraction. A central idea of the IIRA is the need to network larger systems and establish control over hierarchies of machines. Therefore this architecture seems well suited for Industrial Control Systems (ICS). Accordingly, typical IIoT systems are decomposed into five functional, interconnected, domains (control, operations, information, application, business). The IIRA is meant as an iterative top-down process to describe and develop architectural concerns on each viewpoint by their stakeholders. The results of one viewpoint serve as a guideline and impose requirements on the viewpoint below, while discussing the concerns in a subsequent viewpoint may validate or cause a revision of the decisions in the viewpoint(s) above. System architects may use and extend the architectural results of the implementation viewpoint as a basis for a concrete system architecture.

The IIC also published an implementation example for an electrical microgrid testbed that confirms the emphasis on the viewpoints as main tool (Murphy et al., 2016). Furthermore, they produced a joined

paper with the experts of RAMI4.0 to create a mapping between both reference architectures (S.-W. Lin et al., 2017). The researchers highlight the need for interoperability of CPPS and conclude that models, concepts and methods of the reference architectures can be mapped very well, even though the IIRA has a broad cross-industry focus and RAMI4.0 concerns mainly manufacturing.

5.3.3 5C Architecture

The 5C architecture is introduced by Lee et al. and focuses on CPPS manufacturing systems (Lee et al., 2015). 5C stands for the five levels depicted in Figure 26: connection, conversion, cyber, cognition and configuration.

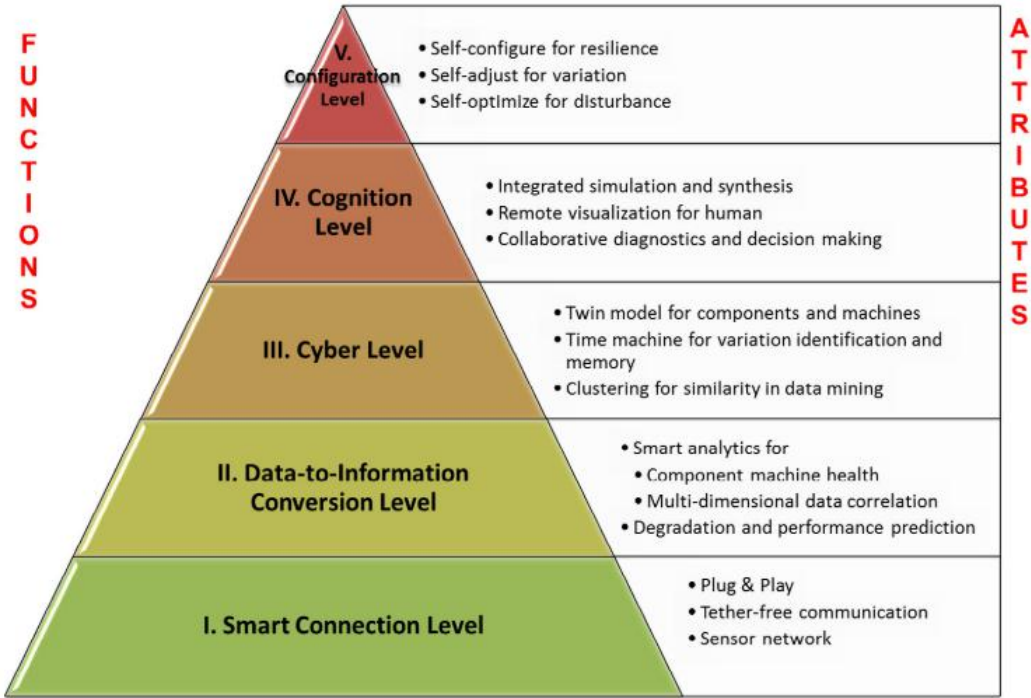


Figure 26 - 5C Architecture

The architecture proposes a guideline for the design of CPPS and how to reach the goal of cognition (here called self-x) starting by the initial data acquisition. So, this architecture aims to provide cognitive functions to the CPS. On the smart connection level, data from versatile sources, such as sensors, controllers, MES, or ERP, are acquired by a central server. Moreover, features, such as sensor signals, are selected on this level. The data-to-information level uses algorithms to process the data and generate information, such as health values or remaining useful life. So typically, machine learning techniques are implemented on this level. The information from many machines is gathered at the cyber level, acting as a central information hub. The digital representation compares a machine against historical data or other machines from the same fleet to get more insights and make predictions. The cognition level provides more in-depth insights to the user, which machine is faulty and what is the origin of the error. In the initial publication, the cognition level appears to be mostly a human-machine-interface, whereas, in later publications, decision-making is the focus of this layer (Bagheri et al., 2015; Lee et al., 2017). It can be used e.g., to optimize the maintenance by prioritizing the tasks or logistic

planning. Configuration level is the top level, which provides feedback from the cyber level to the physical space and implements self-configuration and self-adaption.

The 5C architecture is a more concrete architecture, which aids the implementation of a CPPS in a I4.0 use case. The architecture has a more narrow focus on software development for manufacturing systems, without horizontal integration of suppliers and customers. This leads to Jiang proposing an extension of the reference architecture with three additional facets, i.e., coalition for value chain integration, content for product life cycle information and customer, which addresses individual design, production and after-sales services (Jiang, 2018).

5.3.4 SOAR

The development of Soar was motivated by exploring the requirements for general intelligence, based on the human cognition theory. The first version was created in 1982 as Soar 1 by John E. Laird, Allen Newell and Paul S. Rosenbloom, but the roots of Soar go back to 1956 (Laird et al., 1987). The cognition theory evolved over time up to the current version Soar 9, which is available for free and still an active field of research (Laird & Mohan, 2018).

Soar is a problem-solving architecture and became quite complex since many new features, such as different learning mechanisms, have been added. This introduction focuses only on states and operators as the most basic components of the architecture.

Lehman, Laird and Rosenbloom created the overview for Soar shown in Figure 27 (Lehman et al., 2006). Soar derives decisions via the working memory, which contains the perceptions, the hierarchy of states and the associated operators. To perform a task, Soar uses the initial state, which includes information about the environment as sensed through the perception module, information of current goals and the problem space, and applies operators to change the state within the problem space to reach the destination state.

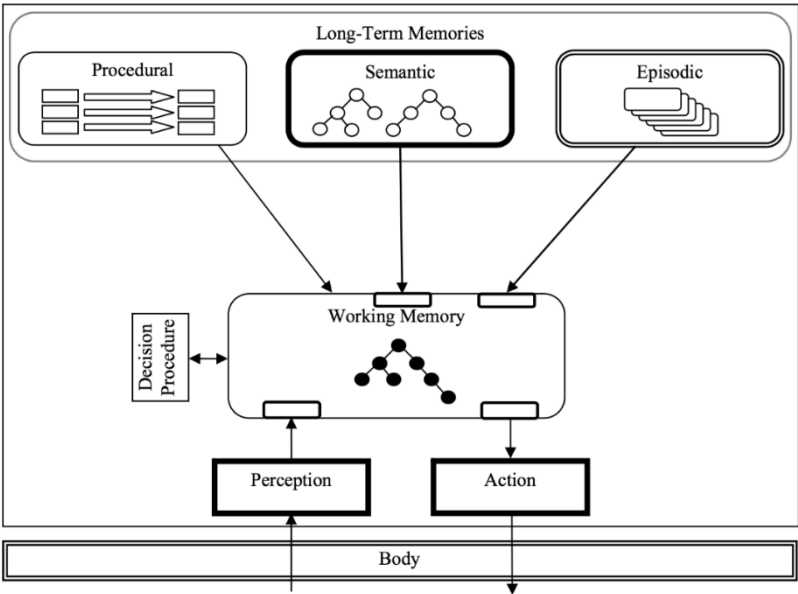


Figure 27 - Soar architecture overview

Each decision cycle evaluates the available operators against the production rules in the working memory and, if necessary, retrieves additional information from the long-term memory, which is divided into three parts:

- Procedural memory encodes rules to go from one state to another
- Episodic memory stores experiences, a history of events, actions and outcomes to predict the results of future actions
- Semantic memory contains knowledge and definitions

Soar uses this knowledge to find all feasible options and select the preferred solution. If there is no feasible operator or more than one optimal solution Soar creates a sub-goal to solve this “Impasse”. Impasses are caused by a lack of knowledge. If an impasse is solved once, chunking is used to preserve the knowledge by creating new production rules, to learn from the situation and not run into the same impasse again. The resolution of the Impasse leads to the selection of a distinct operator, which is applied to the state. This sequence is repeated until the destination state has been reached.

The working memory then translates the decisions into actions, which influence the environment and conclude the decision cycle.

Soar was successfully used in many different use cases, e.g., expert systems for the configuration of computer systems, simulations of air combat through the U.S. Air Force, speech processing and speech creation as well as artificial intelligence for several games (Lehman et al., 2006).

The use in a CPPS seems like a natural choice, as the architecture also possesses sensors and actuators to perceive and influence the environment. The CPPS can use the information about the environment and the production process to derive decisions and implement them. The biggest advantage of Soar is however also the biggest disadvantage; Soar is very generalizable and can model everything, but the implementation effort may be high.

5.3.5 ACT-R

“ACT-R is a cognitive architecture: a theory about how human cognition works. Its constructs reflect assumptions about human cognition which are based on numerous facts derived from psychology experiments” (Bothell, 2019). Anderson’s fundamental idea behind this theory is that, to reach cognition and intelligence “the whole is no more than the sum of its parts, but it has a lot of parts” (Anderson, 1996). Anderson states, that intelligence is the result of gathering and tuning many small units of knowledge. The interaction of procedural and declarative knowledge then leads to cognition.

Procedural knowledge is represented in units called production rules, which mainly consists of a goal, actions and conditions. Declarative knowledge is represented in terms of chunks.

An overview of the ACT-R architecture based on Ritter et al. is shown in Figure 28 (Ritter et al., 2019). The modules are based on the structure of the brain and carry out similar tasks:

- Visual and aural modules process inputs from the external environment.
- The manual module, or motor module, governs hands' movement and translates decisions into the real-world.
- The temporal module provides plausibility checks in the model using time delays (Bothell, 2019).
- The Intentional module, also known as the goal module, maintains the cognitive model and knowledge of the goal.
- The declarative module controls the production and collection of declarative knowledge (Ritter et al., 2019).
- The production system consists of a set of production rules, which represent procedural knowledge and are used to determine the next actions (Salvucci, 2006).

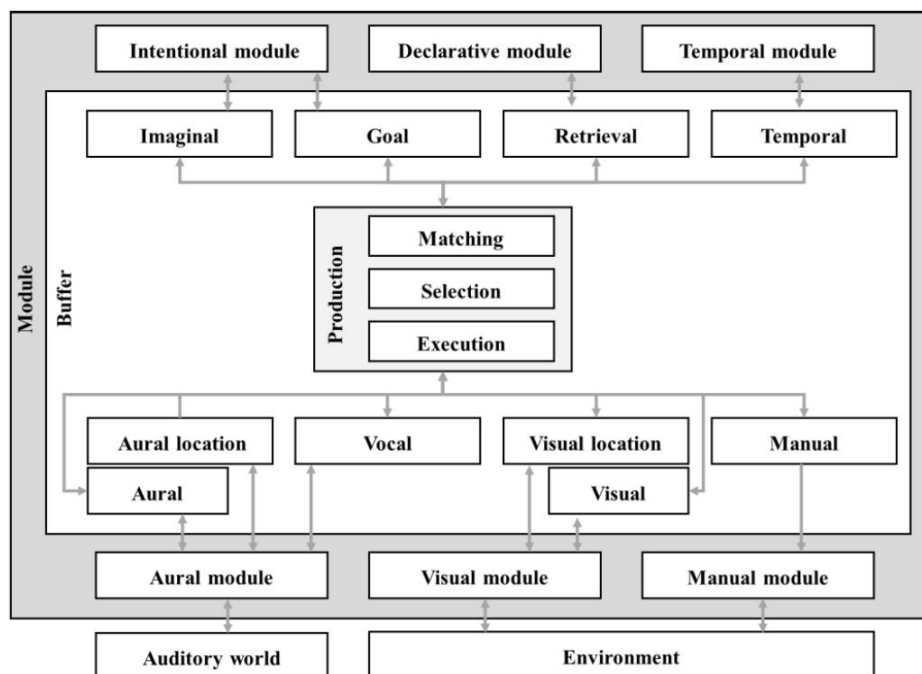


Figure 28 - ACT-R architecture overview

The central production system coordinates the communication and performance of these modules, as they do not communicate with each other directly, through the application of production rules.

The architecture has been successfully implemented for several use cases, e.g., simulating the operation of a car and a plane (Salvucci, 2006; Somers & West, 2012), but also to teach teamwork to autonomous agents in games (Best & Lebiere, 2003). Stewart and West conclude that “ACT-R is a comprehensive theory of human cognition, including how cognition interacts with other mental processes and with the environment” (Stewart & West, 2006). ACT-R’s strength is modeling memory and learning from experiences, similar to Soar, but the implementation effort for a CPPS may also be quite high.

5.4 EVALUATION

The overview of the different candidates gives a first impression for each architecture. The subsequent evaluation uses the consolidated requirements from 5.2.5 to highlight the strengths and weaknesses of the architectures. The consolidated requirements consist of use case specific and general requirements. This section first presents a broad overview of the results and then more detailed results for each of the reviewed architectures.

Table 5 presents the evaluation results. The rows represent the requirements, whereas the columns display the results for the five reviewed architectures. A “-” indicates that the requirement is not addressed by the architecture, while “O” means that an architecture partially fulfills a requirement, e.g., there may be additional effort required. Finally, if an architecture fulfills a requirement, the “✓” is used. The evaluation regards the presented use cases and may differ if the context or the use cases change.

| Requirement | Description | RAMI 4.0 | IIRA | 5C | Soar | ACT-R |
|-------------|--|----------|------|----|------|-------|
| R.1 | Receive declarative user goals | - | - | - | ✓ | ✓ |
| R.2 | Specified interfaces are well defined | o | o | - | - | - |
| R.3 | Strategies to select a suitable algorithm | - | - | - | o | o |
| R.4 | The system learns from experiences | - | - | - | ✓ | ✓ |
| R.5 | Knowledge representation including expert knowledge | ✓ | - | - | ✓ | ✓ |
| R.6 | Acquire data from heterogeneous and distributed systems via versatile protocols | ✓ | ✓ | ✓ | - | - |
| R.7 | Store and manage process data and models | ✓ | ✓ | o | - | - |
| R.8 | Perform preprocessing | ✓ | ✓ | o | - | - |
| R.9 | Learn a model from data with limited time and resources | o | ✓ | ✓ | - | - |
| R.10 | Perform model analysis or simulation with a limited response time | o | ✓ | ✓ | - | - |
| R.11 | Interaction with the user, i.e., human-machine-interface | ✓ | o | - | - | - |
| R.12 | Decision making derives actions, e.g., decide to send new control parameters to the CPPS | o | ✓ | ✓ | ✓ | ✓ |
| R.13 | Apply the action on the controller | - | ✓ | ✓ | ✓ | ✓ |

Table 5 - Evaluation of the requirements for proposed reference architectures

First of all, no reviewed architecture fulfills all requirements. The automation architectures fit generally well, even though several requirements are unfulfilled. Some requirements, e.g., R2, are not adequately fulfilled by any architecture. Just RAMI4.0 and IIRA cover this issue very roughly. Requirement R.3 is also only partially fulfilled by two architectures and depends on the inserted model.

Interestingly, the automation architectures fulfill similar requirements, and the cognitive architectures fulfill similar requirements as well, but those are entirely different from each other. There is a gap in the cognitive aspects of automation architectures. However, cognitive architectures today are not suitable to fill this gap without lacking elsewhere.

The evaluation indicates that the cognitive architectures are best suited to fulfill the general requirements towards a cognitive architecture for I4.0, and the automation architectures better fulfill the use case specific requirements. However, there are likely more requirements towards a cognitive architecture for the implementation of CPPS. The following sections explain the results for each architecture in more detail.

5.4.1 Reference Architecture Model Industrie 4.0 (RAMI4.0)

RAMI4.0 implements all the derived use case requirements on the Layers-axis of the model. The following descriptions concerning the different layers stem from the official RAMI4.0 publication (Adolphs et al., 2015).

The Communication Layer “unifies communication by using a consistent data format towards the Information Layer”. Data transfer is allowed within a layer or between adjacent layers. The Information Layer then “serves structured data via interfaces”. There is no other description or definition for other interfaces between layers (R.2). It is also not possible for the architecture to receive declarative goals from the user (R.1) or for the system to use a learned strategy to select the best algorithm for a problem (R.3) as the system does not learn from experiences in a cognitive sense (R.4). Knowledge representation (R.5) is located in the Business Layer. There the business models, legal and regulatory conditions, and general rules that the system has to obey are represented. The Integration Layer provisions asset information and generates events (R.6). The Information Layer persists data that represents the models (R.7). Pre-processing is also done by the Information Layer, which works as a runtime environment and accepts and transforms data for the Functional Layer (R.8).

The high grade of abstraction makes it hard to use the architecture as a guideline to implement a software system for the described use cases. Those require model learning and analysis, simulation with new data, and decision making (R.9, R.10, R.12). The Functional Layer contains’ runtime and modelling environments for services that support business processes’, but remains unclear about actual functionality implementation. The Information Layer implements’ interaction with the user via human-machine-interface’ (R.11). RAMI4.0 does not intend to lead derived actions back to the controller (R.13). Temporary remote access is allowed for maintenance purposes but not for regular functional integration. Furthermore, the I4.0 component specification concludes that passive communication capabilities are sufficient, i.e., providing status information through a retrofitted higher-level system.

Therefore RAMI4.0 unifies the perspectives of various fields in the different dimensions. That may be simultaneously the biggest advantage and disadvantage. RAMI4.0 can classify machines and components of a CPPS and provide a standardized approach for I4.0 machines’ design. The additional work specifies which data needs to be stored and tries to standardize communication and access to functionality (Federal Ministry for Economic Affairs and Energy (BMWi), 2019; Plattform Industrie 4.0, 2016). However, the architecture may be too abstract for the design and implementation of a CPPS software system.

5.4.2 Industrial Internet Reference Architecture (IIRA)

The IIRA represents a template and methodology that guides system architects to design CPPS systems based on a common framework. However, it quickly becomes apparent that the IIRA does not take all requirements into account. Neither all requirements from the use cases nor the general aspects, as shown in Table 5. The IIRA does not explicitly address declarative goals (R.1). A system architect needs

to further refine the IIRA results to create accurate interface definitions (R.2). Cognitive aspects are not explicitly addressed (R.3, R.4). Expert knowledge seems to be a responsibility of the business domain in the functional viewpoint. However, the architecture provides no further specification of the knowledge representation. The use case specific requirements (R.6-R.13) are in the focus of the IIRA. The information domain in the functional viewpoint regards ingesting sensor data and data quality processing (R.6-R.8) as well as storage and distribution of the data, modeling the state of the CPPS, and performing simulation and optimization tasks (R.9, R.10). Human-machine interfaces (R.11) are located in the business domain in the functional viewpoint. The control domain is also located in the functional viewpoint and addresses the connection of sensors, controllers, actuators, gateways, and other edge systems (R.12, R.13). In conclusion, the IIRA is an architecture that provides high generalizability but is too abstract for a concrete implementation in a CPPS.

5.4.3 5C-Architecture

The 5C architecture stands out with a low level of abstraction even if other architectures fulfill more requirements. The 5C architecture can acquire data, process it, learn from experiences, and adapt the machines, as shown in Table 5. Nevertheless, some requirements are not or not adequately fulfilled. One main issue is the missing HMI of the 5C, so it neither captures declarative goals from the user (R.1) nor informs the user (R.11). The other main issue is that interfaces are not appropriately defined (R.2). There is no explanation of any interface or the transmitted data. The cognitive aspect, to learn from experiences (R.4), is limited to case-based reasoning. Therefore, the architecture compares the current situation to similar situations, e.g., to estimate time to failures, which does not address R.5. The selection of a suitable algorithm (R.3) is another cognitive aspect that the 5C architecture does not address. Some minor aspects are that it is not possible to integrate expert knowledge (R.5). Process data can be lost as just computed data gets persisted, and the preprocessing is not a separate step. However, since an algorithm can integrate the preprocessing, it is indicated as generally achievable in Table 5. Overall the 5C architecture is the most concrete of all reviewed architectures and should enable efficient implementation of CPPS software systems.

5.4.4 Soar and ACT-R

ACT-R and Soar are evaluated together due to the similarity of the architectures. A fundamental feature is the ability to state goals in a declarative manner (R.1). Although the architectures define their interfaces thoroughly, the modular extendability (R-2) is not addressed by the cognitive architectures. There are no implemented strategies for algorithm selection (R-3). However, this could be implemented as a cognitive task and learned by the system. Thus, this requirement is achievable with workarounds. R-4 can be seen as fulfilled, as this is the central task of both architectures. Both architectures' cognition enables decision-making to send parameters to the controller and apply actions in the CPPS (R-12, R-13). Unfortunately, the architectures do not address most use case specific requirements (R-6 to R-11). The low generalizability increases the implementation effort for I4.0 applications dramatically. Thus, the architectures may be less favorable for CPPS development.

5.5 CONCLUSION

None of the introduced architectures fulfills all requirements of a cognitive architecture, as presented in Table 5. Figure 29 classifies the architectures regarding their level of abstraction and generalizability. A high level of abstraction indicates that the implementation has no boundaries, so there is much freedom in implementation. A high degree of abstraction implies that there are no strict guidelines for a concrete implementation. That leads to less guidance for the users, but also greater freedom. In contrast, high generalizability means that the user can efficiently adapt the architecture for many different use cases.




Figure 29 - Gap in proposed reference architectures

The cognitive architectures have a low level of abstraction because a fixed set of structures is available to implement a use case. Although the architectures are cognitive, the generalizability is low, and the architecture has to be adapted for different use cases since each use case requires different knowledge. In comparison to that, the automation architectures are quite abstract. There are no strong guidelines for their implementation, which makes it challenging for the user. However, the reviewed automation architectures can be used in many use cases because they are very generic and possess high generalizability. The 5C architecture is less generic and less abstract than the other automation architectures, but it is still too abstract to intuitively implement it in the use cases. An architecture with a low level of abstraction and high generalizability is required to fill the resulting gap. That architecture is easy to implement and fits many CPPS use cases. Combining the two architecture classes leads to the desired properties. Additionally, the interfaces have to be well defined to enable modularity and reuse existing parts. This need is also mentioned for RAMI4.0 and for the IIRA, where the implementation viewpoint specifies the interfaces (Adolphs et al., 2015; S.-W. Lin et al., 2019). The 5C architecture is a good starting point but needs further concretization.

Thus, an entirely new reference architecture will be presented that tries to fill the existing gap and satisfy the requirements towards a reference architecture in the I4.0 context.

6 CAAI - A BIG DATA REFERENCE ARCHITECTURE FOR I4.0

CAAI – the novel cognitive architecture for AI in CPPS aims to reduce the implementation effort for the usage of AI algorithms in I4.0 applications. The design of the architecture pursues the following four goals, i.e., reliability, flexibility, generalizability and adaptability. These goals ensure continuous production and a standardized implementation for various use cases. CAAI processes declarative goals and learns over time to reach cognitive capabilities and transform the CPPS into an autonomous system. The core of CAAI is a cognitive module. The cognitive module processes the user's declarative goals, selects suitable models and algorithms, and creates a configuration to execute a processing pipeline on the BDP. Constant observation and evaluation against performance criteria assess the performance of pipelines for many and varying use cases. Based on these evaluations, the pipelines are automatically adapted if necessary. The modular design with well-defined interfaces enables the reusability and extensibility of pipeline components and the BDP implements this modular design. The platform uses Docker, Kubernetes, and Kafka for virtualization and orchestration of the individual components and their communication.

Our workgroup developed CAAI as a joined effort and published the resulting architecture in the International Journal of Advanced Manufacturing Technology (Fischbach et al., 2020). The remainder of this chapter is based on and extending that work.

The CAAI BDP, depicted dark grey in Figure 30, wraps the architecture and manages the bus systems and layers. The two layers, i.e., the Data Processing Layer (DPL) and the Conceptual Layer (CL), are shaded in light grey and contain the different processing modules. The three bus systems, i.e., Data Bus, Analytics Bus, and Knowledge Bus, are colored in blue and establish communication between the modules. The arrows demonstrate the modules designated information flow. The cognition module, which establishes automatic configuration, is colored in turquoise. Oval shapes depict external systems, e.g., a CPPS or a HMI.

Data from a CPPS enters the system at the very bottom. The protocol translation module transforms incoming data and sends it to the data bus. The pre-processing module receives raw data, performs the necessary steps to clean the data, and publishes the results back onto the data bus. Other modules in the data processing layer utilize data from the data bus and transfer their analytical results onto the analytics bus. Modules in the CL process information about the user-defined goals and the business logic for a given use case. Those modules evaluate the analytics bus results, determine the parameters to adjust the CPPS via the adaption module, and measure the overall system performance. The CL modules also interact with the knowledge bus to generate reports for the user and process new instructions. The architecture's central element is the cognitive module, which selects and orchestrates different analyses and algorithms, depending on the use case. Therefore, the composition of active modules and their communication over the bus system will change during run time. Providing a pre-defined set of modules and empowering the user to add new modules reduces the overall implementation complexity by building a cohesive yet modular solution.

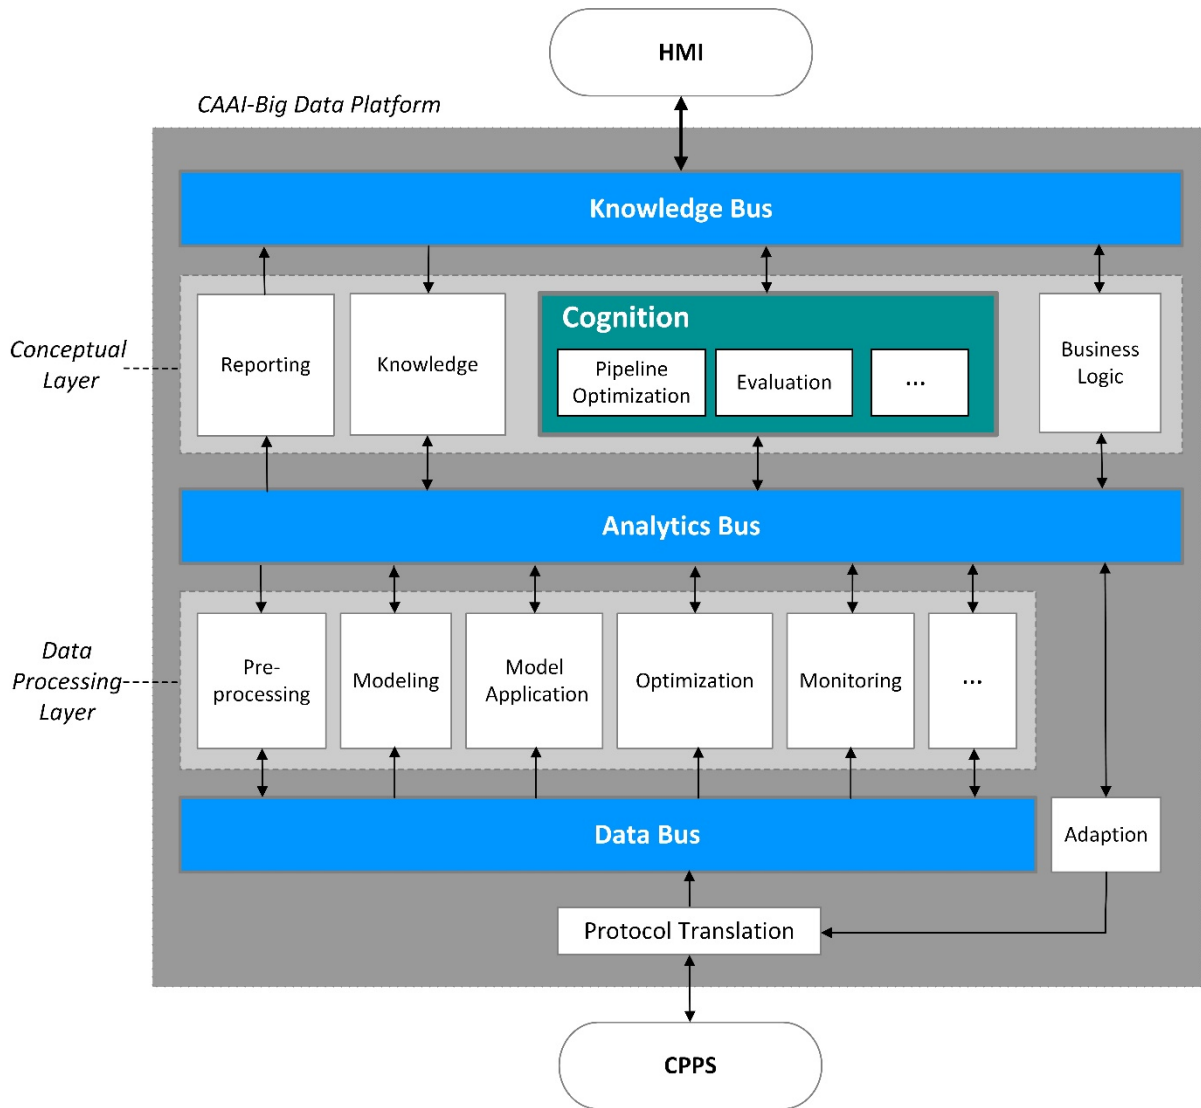


Figure 30 - CAAI architecture overview

The next sections explain the CAAI and the cognitive BDP in more detail. The technologies that are used for the implementation are presented together with the standard building blocks that are provided. The architecture and the building blocks are used to implement three use cases. Finally the architecture is evaluated with the list of consolidated requirements from 5.2.5.

6.1 LAYERS

This subsection introduces the two layers that contain the data processing modules. Both layers encompass individually integrable modules and each module processes the data in a specific manner. Several modules can be combined to reach the overall goal of complex data processing.

6.1.1 Data Processing Layer

The DPL contains all modules that process sub-symbolic data. Pre-processing modules will receive data from the data bus and publish the transformed data back onto the same bus. Those modules that

prepare data for further usage, e.g., by interpolating missing values or synchronizing timestamps, are the minority. Processing modules, e.g., modeling and model application, get data from the data bus and transfer the results to the analytic bus. Modeling modules contain a learning algorithm and can profit from optional expert knowledge. The module publishes the final model to the analytics bus, where model application components can receive and use said model. Again, each of these components has a particular purpose, such as condition monitoring, predictive maintenance, diagnosis, optimization, or similar tasks. The layer also contains a unique component, the adaption. It is unique because it receives commands from the analytics bus, transforms them into a CPPS compatible signal, and sends it directly to the CPPS. This component is adapted to a specific CPPS because it accesses the proprietary CPPS interfaces.

6.1.2 Conceptual Layer

The CL lies between the analytics bus and the knowledge bus. It contains four modules, namely knowledge, reporting, business logic, and cognition. The knowledge module consists of general knowledge and additional information specific to the use case:

- Relevant information about the CPPS, such as signal names, types of devices, or the topology of the CPPS
- General knowledge, such as an algorithm topology, which describes the ability and properties of algorithms
- Constraints that can be defined by the user, e.g., time constraints

The reporting component prepares process data and analytical results for the visualization in the user interface. The business logic verifies that derived actions, such as adapting the CPPS with suitable optimization parameters, do not violate system constraints. The cognition module is responsible for pipeline creation and optimization. The user determines the specific task, and the cognitive module aggregates and configures suitable DPL modules to fulfill that task. Hence, it uses the algorithm topology of the knowledge module as well as past experiences. Monitoring the results of a specific aggregation enables the learning of functional aggregations and improves performance over time. Therefore, the cognitive module is an elementary module of the CAAI and the reason it is called a cognitive architecture.

6.2 BUS STRUCTURE

Different system modules communicate asynchronously via three bus systems, the data bus, analytics bus, and knowledge bus. The processing degree of the data within CAAI grows incrementally from bottom to top and, in some cases, also horizontally. Each bus is a collection of several topics, which can be subscribed by modules attached to the bus. All modules publish their data to the relevant attached bus on pre-defined topics. One or more other modules can subsequently use the data and intermediate results for further processing. Thus, the architecture implements a message-driven processing approach, leading to a flexible and agile system with clear interfaces and hierarchies. The main features of message-driven applications are (IBM, 2016):

- No direct connections between modules,
- communication between modules is time-independent,
- small, self-contained modules carry out work,
- events drive communication,
- data integrity through validation schemata, and

- recovery support.

CAAI aggregates message topics into three bus systems, each with individual properties that are shown in Table 6 and explained in more detail below.

| Property | Data Bus | Analytics Bus | Knowledge Bus |
|----------------------|-----------------------|---------------|---------------|
| Type of data | Raw and pre-processed | Processed | Enriched |
| Velocity | High | Moderate | Low |
| Entropy | Low | Moderate | High |
| Interpretability | Complicated | Moderate | Easy |
| Computational effort | Low | High | Moderate |

Table 6 - Comparison of CAAI Bus Systems

The Data Bus receives the raw production data from a CPPS, demonstrator or simulation. Pre-processing modules clean the data, extract features and publish it back onto the Data Bus as input for other modules from the DPL. The data velocity is high as hundreds of sensor measurements describe a single production cycle. However, the entropy, the information content of all those incoming messages, is quite small, e.g., a temperature sensor sends a measurement every 100ms but the temperature remains constant. As a result of high velocity and a small entropy the interpretability of the production data is complicated and further processing is required. The computational effort for the processing of messages on the Data Bus is low, as preprocessing does not require massive resources.

The Analytics Bus contains the results from the DPL modules and derived actions from modules of the CL, e.g., models with the associated metadata or the Model Application results and adjusted production parameters for the Adaption. Thus, the velocity is moderate in comparison to the Data Bus as hundreds of sensor measurements have been used as input to train and evaluate a few machine learning models. The information content on the Analytics Bus will be higher than on the Data Bus as incoming sensor data has been aggregated and used as input for further processing. The interpretability is also moderate as results have been derived from the vast amount of incoming data. Consequently, the modules that publish results onto the Analytics Bus are expected to use a significant amount of the overall system processing power through training and evaluation of machine learning models.

The Knowledge Bus enables the communication between the user and the system through HMIs and the various modules in the CL, e.g., reporting, knowledge, cognition and business logic. Thus, the bus contains enriched information that is most valuable to the user but the actual amount of data is quite low, e.g., the reporting module aggregates the production data into a single plot to display via the HMI for the user or the user sends a new declarative goal into the system. The entropy and interpretability are high as the Knowledge Bus contains information in the most condensed and enriched form. The computational effort is moderate as the related modules work on previously determined results.

6.3 BIG DATA PLATFORM

The BDP combines several big data applications and utilities. Thus, the BDP supports the following systematic goals:

- (near) real-time data processing
- configuring the system and the goals declaratively
- providing high availability and fault-tolerance
- enable scalability and modularity between different modules

The main tasks are the organization, deployment, operation, and management of big data applications, and the required infrastructure. Infrastructure comprises hosting databases and server processes, such as communication via several bus systems between the system's components and the environment. Therefore, it is possible to integrate structured/unstructured data as well as streaming data. After an initial pre-processing on the incoming data, the big data storage and management will add metadata and persist the cleaned data. The cleaned data is also transferred to the connected bus system to provide the attached modules with new input. Receiving modules can process the incoming data either as a batch or a stream and therefore analyze historic data and continuous streams of live process data. As a cognitive module of the BDP executes and orchestrates different analyses and algorithms, the composition of active modules and the bus system structure will change at run-time. A pre-defined set of modules and the capability to add new modules reduces the system's overall complexity by building a cohesive yet modular solution. The BDP can be hosted on a local server cluster if the hardware is available or data security is especially important. It can just as readily be hosted in the cloud to benefit from flexible hardware provisioning and scalability.

The cognitive BDP utilizes the following concepts and technologies to accomplish the tasks.

6.3.1 Container Virtualization

All components of the system exist as virtualized containers on the CAAI BDP. Isolating the module requirements from the general environment ensures that all requirements for a specific module are met and do not interfere with other modules on the same platform, similar to virtual machines.

In contrast to virtual machines, a container uses the host operating system, and containers share binaries and libraries if possible, which both results in less overhead, as shown in Figure 31 created by Carlos Ferreira from IBM (C. Ferreira, 2015). The small footprint of an individual container allows to dynamically instantiate a large number of containers and assemble modular machine learning pipelines to conduct experiments on the big data platform.

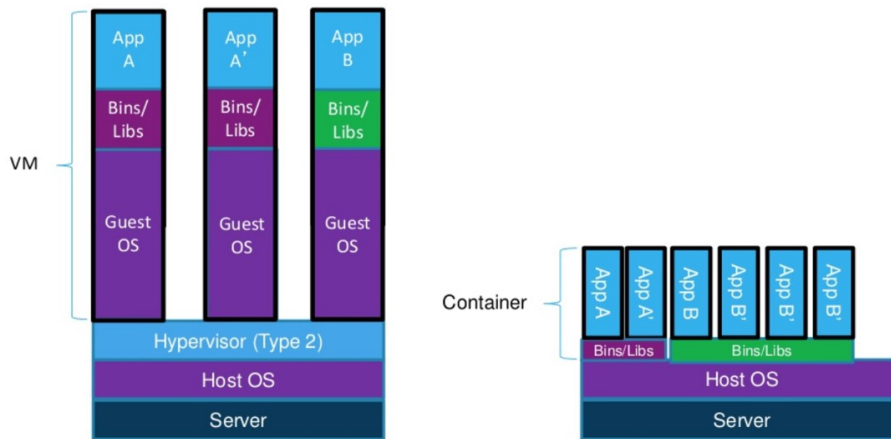


Figure 31 - Comparison: Virtual Machine and Container

The BDP uses Docker²⁰ as container engine. Figure 32 shows an overview of the Docker architecture, which consists of the client, the Docker host and the registry (Docker, 2020). The client communicates via an Application Programming Interface (API) with the Docker host and can use several commands to instruct the Docker daemon. The client can build images based on a Dockerfile, which specifies the base image and all additional requirements to create a specific environment and configuration so the desired algorithm or software can be executed. It is also possible to pull an existing image either from the official Docker registry or another private registry. As every Docker user is allowed to publish images on the central container registry, the Docker registry can be used to provide access to images. Subsequently the client can also run the previously built or downloaded image, which instructs the Docker host to instantiate a container based on this image.

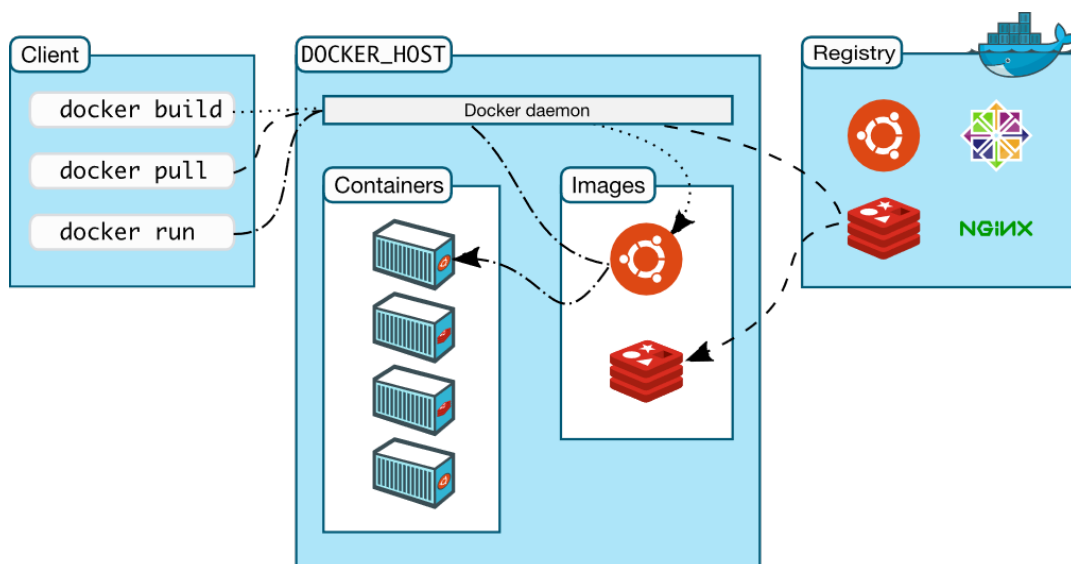


Figure 32 - Docker Architecture Overview

²⁰ Docker can be obtained from the following website: <https://www.docker.com/>

Containers are consistent and immutable, which ensures compatibility across systems. Generally speaking, a validated, running image guarantees to work the same on every computer, server, or cloud-environment. James Turnbull wrote a comprehensive book on the technical details of Docker (Turnbull, 2014).

The prepared images for the CAAI BDP and all components can be found in the KOARCH Docker registry²¹, which enables easy access for potential users.

6.3.2 Orchestration

The CAAI BDP manages the necessary infrastructure and orchestrates virtualized components to compose a system consisting of Microservices that perform a specific task. Orchestration frameworks handle deployments, configuration, updating, and removing of the virtualized software components. A text file declaratively composes a system and lists the different services. Orchestration is done by Kubernetes²², which can utilize the Docker container engine.

An overview of the Kubernetes architecture by Jef Speleta is shown in Figure 33 (Speleta, 2019). A Kubernetes cluster consists of master nodes, which manage the cluster, and worker nodes that run workloads. The user can interact with the master node through the API server either via a graphical user interface or command line tools, such as kubectl²³ or k9s²⁴. Thus, all communication with the cluster happens programmatically, which enables our Cognition to instantiate experiments on the BDP. The scheduling component will then assign the tasks to the worker nodes, while the controller manager oversees the cluster state and takes action if the desired state differs from the actual state, e.g., if a node fails the controller manager will re-distribute the tasks on the remaining nodes. Information about the cluster state is persisted in etcd²⁵, a distributed key-value store.

The worker nodes contain two additional modules, the kubelet and the kube-proxy. The kube-proxy establishes network communication between the master and worker nodes and routes incoming traffic to the correct containers. The kubelet monitors the running containers and sends heartbeat messages to the master node. If the controller does not receive the heartbeat messages measures are taken to restore the cluster state as the worker node is supposed to malfunction.

²¹ KOARCH Docker Hub Profile: <https://hub.docker.com/u/koarch>

²² Kubernetes can be obtained from the following website: <https://kubernetes.io/>

²³ Kubectl can be retrieved from the following website: <https://kubernetes.io/docs/tasks/tools/install-kubectl/>

²⁴ K9s can be obtained from the following website: <https://k9scli.io/>

²⁵ Etcd Project Website: <https://etcd.io/> retrieved 28.11.2020

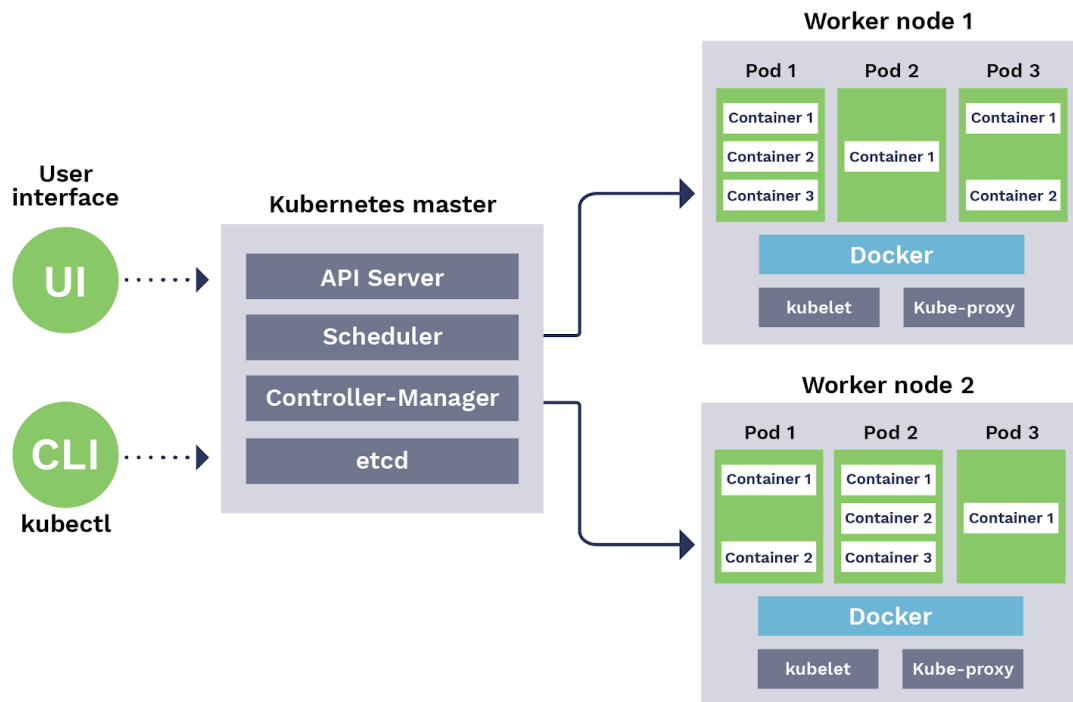


Figure 33 - Kubernetes Architecture Overview

Thus, Kubernetes assists the continuous operation of the big data platform. Furthermore, it allows the Cognition module to use the orchestration to instantiate pipelines with selected algorithms and evaluate the results.

More details on the technical implementation of Kubernetes can be found in the publication by Hightower, Burns and Beda (Hightower et al., 2019).

6.3.3 Microservices

All modules are developed as Microservices to compose the software system for a specific use case from smaller self-sufficient parts. Each module includes standardized communication functionality to publish and subscribe to relevant topics on the bus system. The resulting system is modular, language-agnostic, and utilizes well-defined interfaces. According to Microservice best-practices, each Microservice can store internal data in its local storage (Watts, 2015).

The comparison with a traditional monolithic application, shown in Figure 34, highlights the benefits of Microservices. The monolith is a single big application, e.g. a machine learning system, deployed on a server. It is not possible to run some modules of the software on another server, therefore a lot of processing power and memory are required in a single machine, which is costly. If the requirements of the software change the only option is to scale vertically, which means to upgrade the server or buy a newer machine and move the application. Decomposing a big application into smaller Microservices, e.g. a module that handles the data preprocessing and another module that trains a machine learning model, enables running the Microservices on separate servers. The server cluster can be upgraded by adding another server and distribute the workload evenly, which is called horizontal scaling, if more processing power is needed.

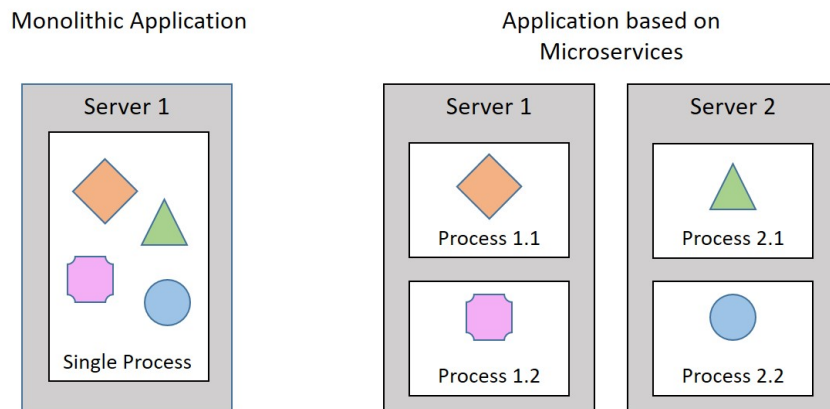


Figure 34 - Monolithic application and Microservices

Composing a software system with Microservices also allows to scale each service independently. Therefore it is possible to adjust the instances of a service according to the workload, as shown in Figure 35. Process 1.1 (orange) and 2.1 (green) are Microservices with a workload that do not require much processing power, e.g., a data preprocessing module. It is sufficient to run a single instance of each Microservice. The other processes (pink and blue) may represent machine learning algorithms, which require a lot of resources. It is possible to execute the algorithms with several instances in parallel, which leads to a faster computation and better resource utilization.




Figure 35 - Scalability of Microservices

Thus, the Microservice approach with horizontal scaling proves to be more cost-effective as it is possible to use cheaper hardware while also using the available resources more efficient (Syafudin et al., 2017). Furthermore, it allows the big data platform to exchange individual modules from a machine learning pipeline, while keeping the overall system in continuous operation.

6.3.4 Messaging

The different bus systems managed by the CAAI BDP transfer data via messaging. Messaging allows asynchronous communication between modules and enables parallelization as well as processing data

several times for different purposes via topics and consumer groups. Kafka²⁶, a project from the Apache Software Foundation, is used to implement a reliable message system for the platform.

Kafka is a distributed system, thus replication of messages is possible via multiple broker instances. If a single node fails, the Producers and Consumers can communicate with the remaining Brokers. Additionally Kafka is a log-based system and all messages sent via Kafka are persisted. Each Consumer is able to decide the starting point for message processing within a topic. Thus it is possible to process messages again and restore past application state, if the whole system should fail. Both mechanisms ensure a very fault-tolerant communication for the Microservices.

The publish-subscribe pattern gives the necessary flexibility to conduct a variable number of experiments in parallel. Figure 36 depicts a producer sending information to a message broker and several Consumers in two Consumer Groups that retrieve and process the messages. The publish-subscribe mechanism works as follows:

- 1) The Producer sends three distinct messages for topic T1 to the Broker. Consumers can subscribe topic T1 to receive its content.
- 2) Consumer Group 1 consists of two Consumers (1+2). The Consumer Group subscribes to a topic to receive new messages from the Broker whenever a Producer publishes messages to this Topic. As the Consumers belong to the same Consumer Group they will work on the same task and process the three messages in parallel, which is useful if a task is very time-consuming or response time is restricted.
- 3) Consumer Group 2 consists of a single Consumer (3). Both Consumer Groups receive all messages and can process them for different purposes and with different grades of parallelism.

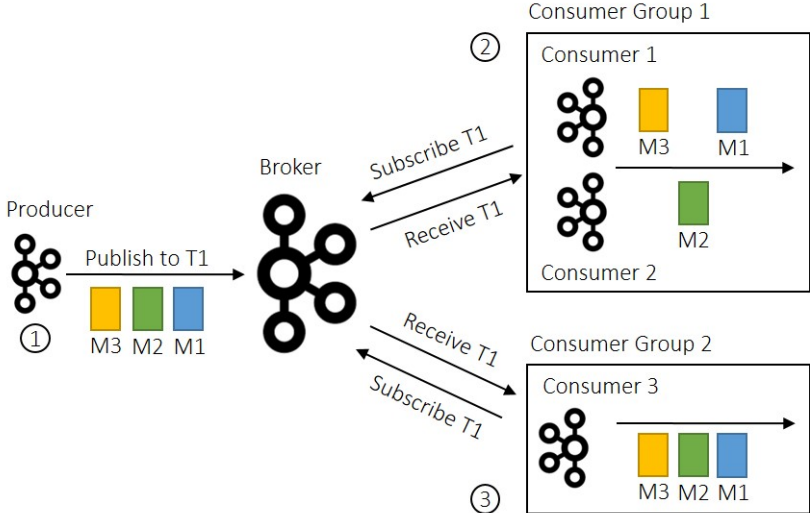


Figure 36 - Messaging via publish / subscribe

Now that general communication between modules of the BDP is established, Consumers and Producers can be used to chain different operations together and use intermediate results as input for

²⁶ <https://kafka.apache.org/>

the next step. The diagram below shows a Producer P1 creating messages and publishing those to a Topic T1. The intermediate module consists of a Consumer C1 that subscribes to Topic T1, transforms the incoming messages and publishes the results to Topic T2 via Producer P2 for further use. Finally a Consumer C2 receives the messages with the intermediate results from Topic T2 and is, for example, able to visualize the output or persist it into a database.

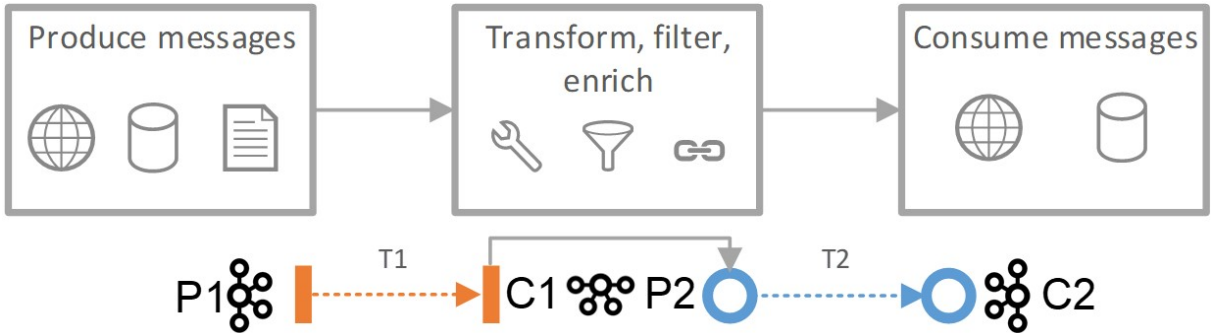


Figure 37 - Chaining Producers and Consumers to create a processing pipeline

Further information on the technical details can be found in the very thorough book from Narkhede, Shapira and Palino (Narkhede et al., 2017). A working example can be found in my Github Account²⁷.

Using Kafka as messaging technology for the BDP enables fast, continuous and fault-tolerant communication between all modules via the Bus system.

6.3.5 Schema Management

A schema stores the metadata of the data, with all the available fields and datatypes (Confluent, 2020). When a module publishes to the bus system, the serializer applies the schema and encodes the message or filters out non-conforming messages. A consumer that subscribes to a topic on the bus has access to the same schema and can verify the integrity of incoming messages. A central schema registry distributes and versions the schemas which allows regulated data evolution. The BDP implements schema management through Avro²⁸, another Apache Software Foundation project. Avro is often used in big data projects because it provides rich schemas, fast serialization and needs less storage space than more traditional JSON or CSV formats (Draken, 2019). The Code Listing 1 below shows an example of an Avro schema.

Lines 1-3 consist of metadata and indicate that the schema defines records for a user in the example namespace. Lines 4-8 define the fields for each record, e.g., the user’s name, favorite number and favorite animal with the associated data types. Line 5 specifies the user’s name as a string with no other conditions or limitations. Line 6 defines the favorite number, but the additional type “null” marks this field as optional. The favorite animal is defined in line 7 and provides the valid choices “cat” and “dog”.

²⁷ Github Account Jan Strohschein: <https://github.com/janstrohschein/KOARCH/tree/master/Big%20Data%20Platform/Docker/Kafka>
²⁸ Avro can be obtained from the following website: <https://avro.apache.org/docs/current/spec.html>

Using schemata to define the fields and datatypes ensures clear communication for the BDP via the bus system and additional modules can be integrated easily as the semantics for each record are provided.

```

1  {"namespace": "example.avro",
2  "type": "record",
3  "name": "user",
4  "fields": [
5    {"name": "name", "type": "string"},
6    {"name": "fav_number", "type": ["null", "int"]},
7    {"name": "fav_animal", "type":
8      {"type": "enum", "name": "animals", "symbols": ["cat", "dog"]}}
9  ]
10 }

```

Code Listing 1 - Avro Schema Example

6.3.6 Technology Summary

The combination of those five technologies enables the big data platform to create a modular machine learning system, shown in Figure 38. Prepared container images for the modules can be obtained from the official CAAI repository and are automatically stored in the local container registry (T1). Kubernetes, as orchestration engine (T2), uses these images to instantiate the other infrastructure modules, e.g. the Message Broker (T4) and the schema registry (T5), to enable a standardized communication on the BDP. Subsequently the Cognition module can instruct the Orchestration to start the Microservices (T3) to compose one or more machine learning pipelines. Those Microservices are autonomous and language-agnostic, which means every module can be written in the language that is most appropriate for the task at hand.




Figure 38 - Cognitive Big Data Platform

The presented technologies are used to implement a set of standard building blocks. Those building blocks are generic enough to be valuable in various use cases and portrayed in the following section.

6.4 STANDARD BUILDING BLOCKS

Standardizing and creating common basic building blocks transfers the architecture into practice. This process is important as it demonstrates design principles for the architecture modules. Providing a set of standardized building blocks also reduces the implementation effort for future use cases as common functionality is already available, e.g., equipping all modules with message functionality, plotting the data in real-time or storing the production data in a database. The following section introduces the proposed module structure, the container build process and several example building blocks that can be used to compose machine learning pipelines. All building blocks with working examples and further instructions can be found in the Github Repository²⁹. The building blocks in the Github Repository all use the same structure to provide an easier orientation. Figure 39 shows an overview of the architecture building block structure.

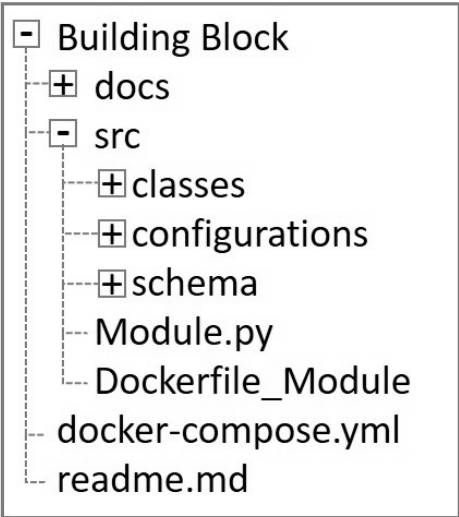


Figure 39 - Building Block Structure

Each building block provides two entry points at the top of the folder structure, i.e., the `readme.md`, which contains further information and instruction for the user, and the `docker-compose.yml` with the technical specification for a building block. The complete structure of a building block and the individual purpose are explained in Table 7.

²⁹ Jan Strohschein Github Repository – Building Blocks:
https://github.com/janstrohschein/KOARCH/tree/master/Big_Data_Platform/Docker
retrieved 15.11.2020

| File/Folder | Purpose |
|----------------------|---|
| └ docs | Contains images, diagrams and other supporting material for the documentation. |
| └ src | The source folder stores the code files as well as the Dockerfiles. Dockerfiles consist of instructions to build a container. The Dockerfile is named after the module it containerizes. |
| └ classes | The subfolder contains the different classes, e.g., the Kafka class, for other modules to inherit. |
| └ configurations | The configurations folder stores several files: <ul style="list-style-type: none"> ▪ The `config.yml` contains the configuration for all modules in the project. Furthermore, the config file includes general use case specific configuration, e.g., regarding the objective function or the initial design. The description of each service in the docker-compose file specifies which sections of the config file will be used in the container. ▪ The requirements.txt contains all the packages that need to be installed in a container. The Dockerfile copies the file into the container and installs the packages during the build process. |
| └ schema | The modules use indirect communication with a messaging approach. All messages are sent via Kafka and verified with the related Avro schema, which is stored in an avsc-file. Each module specifies its input- and output-topics and the associated schemas in the `config.yml`. |
| └ docker-compose.yml | Several services can be combined into a docker-compose file, which allows to manage all services together. Each service entry consists of: <ul style="list-style-type: none"> ▪ The container and host name. ▪ Build information such as the base image or the path to the Dockerfile. ▪ Assign files or volumes from the host system to a specific path in the container file system. ▪ Environment variables, e.g., the config path and the config sections relevant to the module. ▪ Port forwarding from the host system to the container. Furthermore it is also possible to define a common network for all containers and to specify how Docker volumes should be used. |
| └ readme.md | The readme.md explains the module/use case and gives usage instructions. |

Table 7 - Building Block Structure and Purpose

A building block can be started through the `docker-compose.yml` file and the container creation and instantiation process is shown below in Figure 40. The `docker-compose.yml` file contains the description for all services that compose the building block, defines virtual networks for the services and allows to mount parts of the host file system as volumes into a service. Each service contains more detailed information for the build of its individual container, e.g., the relevant Dockerfile and ports that should be forwarded from the host system into the container. The service description also defines which sections in the configuration file `config.yml` are relevant for the specific service. The Dockerfile specifies which image should be used as the foundation for the container. A local image cache stores copies of images and loads the image from the global Docker registry if it does not exist. Thus, the cache accelerates the instantiation of several containers based on the same base image. Docker extends this base image through installation of the required software packages. Finally, it copies the building block source files into the container. Thus, it provides an identical and isolated application environment for each use, independent from other applications on the host system or other versions of the same software.

Each building block can be configured via the `config.yml` file. This file is divided in sections and contains general information for all containers, e.g., the address of the Kafka Broker, and additional service-specific information.

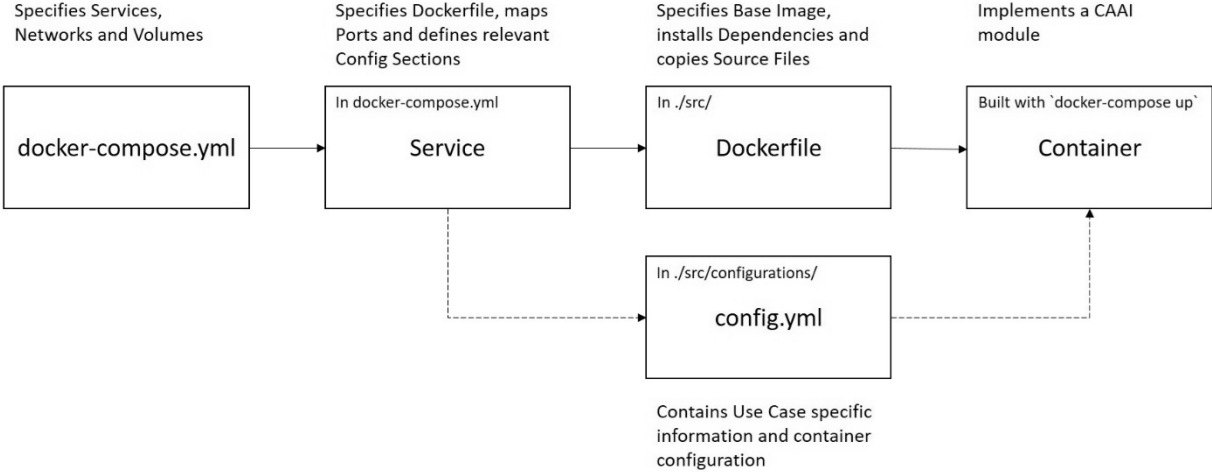


Figure 40 -Instantiating a CAAI Building Block

The following sections present several of the implemented standard building blocks that provide commonly used functionality. These modules are also used in the following section 6.5 to implement the use cases.

6.4.1 Kafka Broker

Kafka consists of brokers, producers and consumers. The Kafka Broker receives and distributes messages and is an integral part of the BDP. A producer publishes a message to a certain topic and sends it to the broker. A consumer subscribes to a topic and receives incoming messages. The necessary Kafka messaging infrastructure, i.e., Zookeeper, Kafka Broker and the Schema Registry, can be started via the corresponding `docker-compose.yml` file. The implementation and further instructions are available in the Github Repository³⁰. Docker loads the base image from Docker Hub and builds the containers for the three services. Figure 41 depicts the start-up and operation of the Kafka messaging system including the following steps:

1. Start the Zookeeper Server. Zookeeper handles the orchestration between multiple Kafka Brokers and saves their application state. If the restart of Kafka is necessary the last state can be restored so that all messages are processed correctly.
2. Start the Kafka Server. In a message-based infrastructure it is possible to have several Kafka Brokers on the same server or Brokers on multiple servers communicate with each other.
3. Start the Schema Registry. The Schema Registry manages the message schemas and enables schema evolution, which means the Schema Registry stores multiple versions for a schema if

³⁰ Jan Strohschein Github Repository - Kafka Broker:
https://github.com/janstrohschein/KOARCH/tree/master/Big_Data_Platform/Docker/Kafka_Broker
 retrieved 24.11.2020

the data schema changes over time. Producers and Consumers can retrieve a specific version of the schema or the most recent version.

4. Instantiate a Producer and connect it to the Kafka Broker and the Schema Registry.
5. The Producer sends a message to the Broker. New messages can be created from various sources, for example the internet, databases, files or also production systems. The Kafka system uses the binary Avro format to encode the messages. Benefits of this approach are message compression and checks of data validity against the schema when it is encoded. The Producer registers the schema for the messages with the Schema Registry. That allows to include the `schema_id` instead of the schema itself when sending new messages, which reduces the communication overhead.
6. Instantiate a Consumer and connect it to the Kafka Broker and the Schema Registry. The Consumer subscribes a specific topic from the Kafka Broker and retrieves the corresponding schema from the Schema Registry.
7. As soon as new messages arrive at the topic the Consumer will collect them and decode the message to restore the values. Finally, the Consumer can process the data and create an intermediate result for further consumption or a final output.

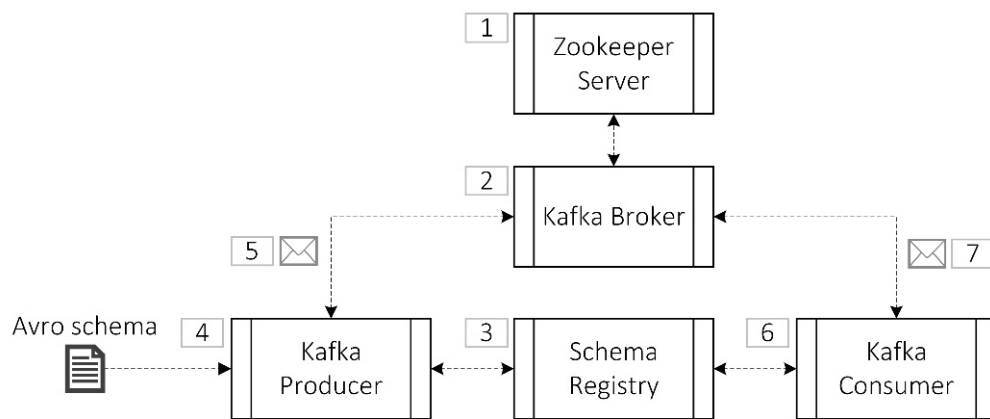


Figure 41 - Kafka Workflow

6.4.2 Kafka Client

The KafkaPC client module is a basic building block that provides message functionality which is necessary for communication between modules on the BDP. Other CAAI modules inherit the functionality to consume and/or produce messages:

- Producer publishes messages to a certain topic (P)
- Consumer transforms intermediate results and produces new messages (PC)
- Consumer receives messages and creates an end-result but no new messages (C)

Thus, the module consists of general functions, e.g., creating a connection to the Schema Registry or reading the supplied configuration file, specific functions for the Producer and Consumer and the provided configuration, as shown in Figure 42.

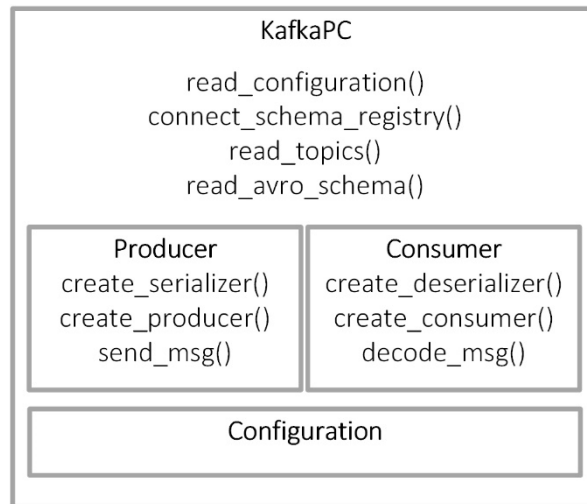


Figure 42 - KafkaPC Client Module

The KafkaPC client module is available in the Github Repository³¹. It can be launched with the corresponding `docker-compose.yml` file, which builds the container and starts the Kafka Clients as follows:

1. The module processes the configuration, retrieves the information for Producers / Consumers and establishes the connection to the Schema Registry.
2. The client creates serializer for outgoing and de-serializer for incoming messages based on the message schemas associated with the topics.
3. The serializer and de-serializer objects are used to instantiate Producers and Consumers based on the configuration and connects those to the Kafka Broker and the Schema Registry.
4. Finally, the KafkaPC Client module can produce and consume messages on the CAAI platform.

6.4.3 Postgres Database

PostgreSQL³², also known as Postgres, is a free and open-source relational database management system with over 30 years of active development that is written in the C programming language. The database building block ensures that other modules are able to persist data, model relationships between entities and query the data for further analysis.

The CAAI container implementation for the PostgresDB building block and a working example can be found on GitHub³³. The example implementation uses two `docker-compose.yml` files, one to build the container and prepare the database, and another for continuous operation.

During the initialization phase Docker builds the PostgresDB container and initializes the PostgresDB volume for persistent data storage, as shown in Figure 43. The PostgresDB container uses the

³¹ Jan Strohschein Github Repository - Kafka Client:
https://github.com/janstrohschein/KOARCH/tree/master/Big_Data_Platform/Docker/Kafka_Client/Confluent_Kafka_Python retrieved 24.11.2020

³² PostgreSQL Project Website: <https://www.postgresql.org/> retrieved 15.11.2020

³³ Jan Strohschein Github Repository - PostgresDB Example:
https://github.com/janstrohschein/KOARCH/tree/master/Big_Data_Platform/Docker/DB/Postgres retrieved 15.11.2020

instructions from the `init.sql` file³⁴ during the initialization phase to create the tables, the fields with the associated data types and the relationships between the tables. This script can be adapted to model the data for the specific use case.




Figure 43 - Postgres Initialization Phase

Once the preparations are complete, the second Docker-Compose file can be used to start continuous operation. Docker builds the Postgres module, which is used to interact with the previously prepared PostgresDB container and thus also the PostgresDB volume, as shown in Figure 44. The Postgres module can be adjusted for the specific use case and decoupling the containers ensures that several Postgres modules can connect to the PostgresDB database at the same time. It is also possible to change and re-build the Postgres module without losing the data that is stored in the actual database.



Figure 44 - Postgres Operation Phase

The Postgres module inherits the message capabilities from the KafkaPC module, as shown in the detailed view in Figure 45. Furthermore, it adds the functionality to open a connection to the Postgres DB container and execute a prepared SQL statement. Thus, it is possible to subscribe to any topic, receive the messages, perform transformations as necessary and persist the data in the database.

³⁴ SQL stands for Structured Query Language, a “language used to access and modify information in a database. It stores SQL statements for creating or modifying database structures, insertions, updates, deletions, or other SQL operations”. For more information see: <https://fileinfo.com/extension/sql>

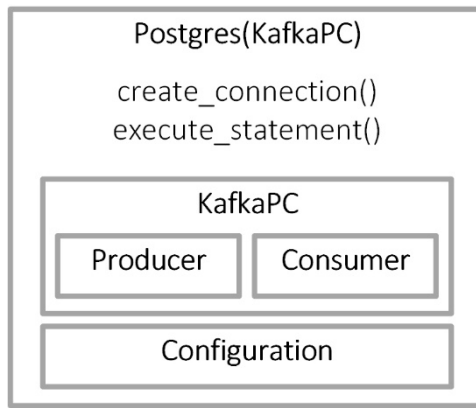


Figure 45 - Postgres Module

The example use case consists of three containers and a volume to establish the database, receive data from arbitrary sources and store it in a persistent volume, as shown in Figure 46 below:

- **L0_P_Send_Users**
The module reads a list of users and sends the messages via Kafka to the specified topic on the Kafka Broker.
- **L1_C_Persist_to_Postgres**
The Postgres module subscribes to the topic and receives the data. It establishes a connection to the PostgresDB container and sends the data that should be persisted.
- **PostgresDB Container + Volume**
The PostgresDB container persists the data in the database. Detaching the PostgresDB container and the physical volume that holds the data allows to persist data independently from the container life-cycle. Thus, no data is lost in case of database failure or container crash.




Figure 46 - Postgres Module persists incoming messages in PostgresDB

Further usage of the Postgres module is shown in the emotion detection use case in 6.5.1. There the database is used to store the profile information of individual users, their status updates and the associated emotional classifications. As PostgresDB is a relational database, it is easy to join the data from the different tables for further analysis.

6.4.4 Plotting

The plotting module can be used to create a real-time dashboard from messages on the BDP. Krystian Schindler developed the first iteration of the module during his bachelor thesis in the KOARCH project (Schindler, 2020). Since then, the module has been collaboratively integrated into the CAAI architecture and has been developed further. The working example of the plotting module is available in the KOARCH Github repository³⁵.

The example plotting application consists of three modules:

- **L0_P_Send_Data**
Produces messages with data from the production process and publishes those to separate topics via Kafka as specified in the `./src/configurations/config.yml``.
- **L1_C_Reporting**
Collects the messages from the designated topics, optionally processes the data with specified functions and sends the messages to the plot topic.
- **L2_C_Plot_Data**
Retrieves the messages from the plot topic and creates the plots in real-time. It is possible to plot data from various sources and the web interface dynamically creates additional tabs for each combination of data source and plotting variable. Incoming data can also be grouped by an additional variable, in this example the utilized algorithm.

The plot module expects messages from L1_C_Reporting to be encoded with the Avro schema, which has been introduced in 6.3.5. The content of the `./src/schema/plot.avsc`` file is shown below:

```
1  {"type": "record",
2   "name": "Plot",
3   "fields": [
4     {"name": "plot", "type": {
5       "type": "enum", "name": "plottypes", "symbols":
6         ["single", "multi"]}},
7     {"name": "multi_filter", "type": ["null", "string"]},
8     {"name": "source", "type": ["string"]},
9     {"name": "x_label", "type": ["string"]},
10    {"name": "x_data", "type": ["int", "string"]},
11    {"name": "x_int_to_date", "type": "boolean"},
12    {"name": "y", "type": {
13      "type": "map", "values": ["float", "string"]}
14  ]
15 }
```

Code Listing 2 - Plot schema

³⁵ Jan Strohschein Github Repository - Plot module:
https://github.com/janstrohschein/KOARCH/tree/master/Big_Data_Platform/Docker/HMI/Plot retrieved 14.11.2020

The different fields in the schema serve the following purpose:

- "plot" is defined as an enumeration, with the available values "single" or "multi", which specifies if a single variable is shown, or if data is grouped by another variable.
- "multi_filter" signifies the variable to group a multi-plot, not defined for a single plot.
- "source" is the name of the data source. Will be used as prefix for the tabs in the web interface.
- "x_label" specifies the label for the x-axis.
- "x_data" assigns the variable that should be displayed on the x-axis. Can be an integer value, e.g., an id, or a string, e.g., a timestamp.
- "x_int_to_date", set this to True if "x_data" is an integer representation of a timestamp to automatically convert the input. Otherwise this is set to False.
- "y" contains all the data that should be plotted. Each entry will create another tab with a new plot. It also needs to contain the data for the multi-filter, if the plot is a multi-plot.

Following the instructions in the Github repository starts the Kafka broker, the plot server and the data processing modules. The web interface leads to the different plots as shown below in Figure 47 and Figure 48. The first plot shows the different values for the production parameter x over 50 production cycles. The color of each data point is determined by the algorithm that was allowed to adjust the parameter. The first five iterations use values from the initial design. The points are equally distributed over the search space for parameter x. The machine learning algorithms use the results from those five initial iterations to train the respective model. After the initial phase each algorithm predicts the best new value for the parameter and the evaluation module chooses the best algorithm. The Kriging algorithm initially got better results and was allowed to set the next values for the production parameter. However, with a few more training samples the Random Forest algorithm could provide better results or equal results with less resource usage and determined the parameter values for most production cycles.

The second plot presents the resource usage of both machine learning algorithms for each production cycle. Random forest uses nearly constant amounts of processing time, while the resource usage for the Kriging algorithm increases with the growing amount of additional data points, as expected.

Those two plots are exemplary for the functionality of the plotting module. The plot module is able to plot any numeric data of interest and can be configured declaratively.



Figure 47 - Plot of the new process parameter value and the deciding algorithm




Figure 48 - Plot of the CPU resource usage per Algorithm

6.4.5 FastAPI

FastAPI is used to implement CAAI APIs according to the OpenAPI specification³⁶. The framework enables a quick and straightforward implementation and automatically generates a web interface based on the API specification.

The basic implementation example, which can be found in the Github repository³⁷, presents a generic API, that creates the necessary endpoints to handle the API requests based on the included configuration file. Dynamically creating the endpoints for the API is important as the composition of the modular system is not fixed and configuring it declaratively prevents additional implementation effort. The example creates a simple data processing pipeline consisting of four modules:

- **L1_P_count & L1_P_double_count**
Produce simple messages and publish those to separate topics via Kafka.
- **L2_C_Reporting**
Collects the messages from all specified topics and sends the messages to different API endpoints with optional processing.
- **L3_API_HMI**
Dynamically creates endpoints based on the configuration and offers the data as JSON or CSV.

The configuration file `./src/configurations/config.yml` determines the API endpoints. The YAML format uses indentation to represent nested groups of key-value pairs. Code Listing 3 shows the relevant excerpt from that file to configure the module `L2_C_Reporting`.

```
1 L2_c_reporting
2   IN_TOPIC:
3     AB_counts: ./schema/count.avsc
4     AB_double_counts: ./schema/count.avsc
5   IN_GROUP: reporting
6   API_OUT:
7     AB_counts: forward_topic
8     AB_double_counts: forward_topic
9   API_URL: http://L3_API_HMI:8000
10  API_ENDPOINT: /topic/
```

Code Listing 3 - FastAPI Configuration

³⁶ “The OpenAPI Specification (OAS) defines a standard, programming language-agnostic interface description for REST APIs, which allows both humans and computers to discover and understand the capabilities of a service without requiring access to source code, additional documentation, or inspection of network traffic.”, for more information see: <http://spec.openapis.org/oas/v3.0.3> retrieved 31.10.2020

³⁷ Jan Strohschein Github Repository - FastAPI Example:
https://github.com/janstrohschein/KOARCH/tree/master/Big_Data_Platform/Docker/HMI/FastAPI retrieved 09-11-2020

The content of the excerpt has the following purpose:

- Line 2-4: Instructs the reporting module to collect the messages from the topics `AB_counts` and `AB_double_counts` and decode the incoming messages with the specified schemas.
- Line 5: Specifies the consumer group and ensures that other consumers can process the same messages for other purposes.
- Line 6-8: Defines how the incoming messages are treated before being sent to the API, e.g., the messages from topic `AB_counts` will be processed by the function `forward_topic` within the reporting module.
- Line 9-10: The reporting module sends the processed data to the API. The final address for the POST request is the combination of API_URL and API_ENDPOINT. The hostname `L3_API_HMI` represents the container name of the API module. Specifying the hostname like this works because Docker creates a virtual network for all containers.

FastAPI provides an automatically generated web interface based on the API description, which presents an overview over the available API routes, as shown in *Figure 49*.



Figure 49 - FastAPI HMI Overview

Each of the four visual blocks represents an API route and the associated HTTP method as defined in the standard building block `L3_API_HMI.py`. The HTTP methods refer to the standard that was formulated by the World Wide Web Consortium (World Wide Web Consortium, 1999). The “GET” method indicates data retrieval from the API, whereas the “POST” method signifies data transfer to the API. The four routes accessible via the web interface serve the following purpose:

- (GET) Return all results as JSON via: /topics/
- (GET) Return results as JSON for a specific topic via: /topic/{topic_name}/
- (GET) Export the results for a specific topic to CSV via: /topic_csv/{topic_name}/
- (POST) Add a row to a specific topic via: /topic/{topic_name}/

The API routes can contain a variable path, i.e., `{topic_name}`, to target the content of specific topics, as specified in Code Listing 3.

The user can select a route to get more information and execute the associated query. Figure 50 shows the interface to add a single data row to the result set of a specific topic. FastAPI utilizes the information directly from the standard building block `L3_API_HMI.py` to provide the information how a specific endpoint can be used. The `topic_name` dropdown-menu contains all the topics that have been specified in the configuration file. Thus, the user can select a topic, enter the information for a new row and execute the query against the API. FastAPI will retrieve and display the server response to confirm that the process was successful. The interface also shows the equivalent Curl request³⁸, so the query can be translated into code, as done in `L1_P_count` and `L1_P_double_count`.

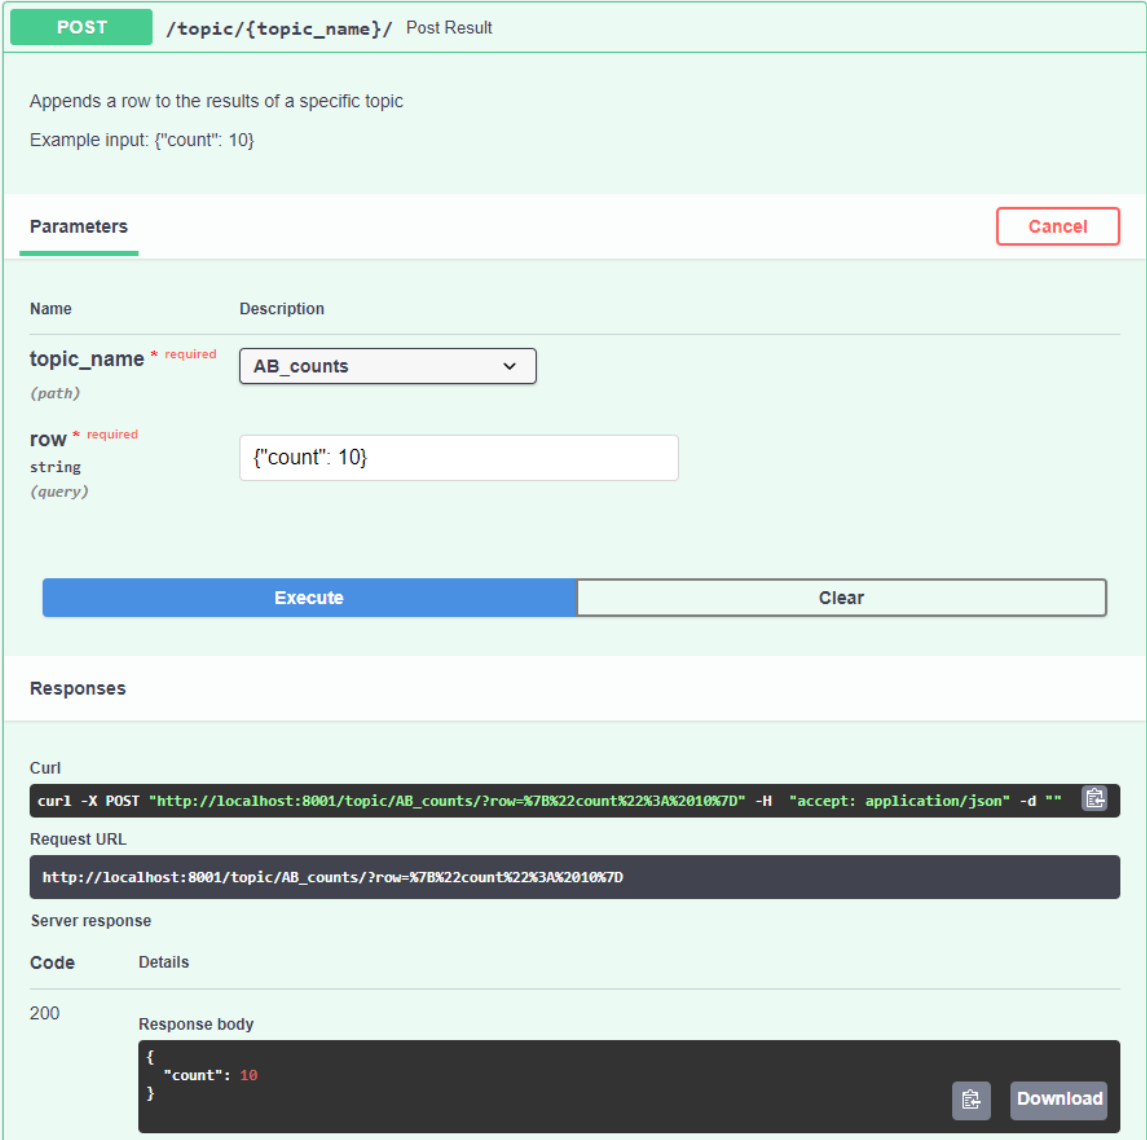


Figure 50 – FastAPI Route Detail

³⁸ “Curl is a command line tool and library for transferring data with URLs”, from the project website: <https://curl.se/> retrieved 09-11-2020

6.4.6 Knowledge API

The knowledge API module stores information about the use case and knowledge about the available algorithms. FastAPI, as presented in the previous section, is used to implement the knowledge building block. The working example for the knowledge building block is available in Github³⁹.

The module allows the user and other modules such as the Cognition to retrieve, update and extend the available knowledge. Code Listing 4 presents the knowledge representation of a Random Forest algorithm for an optimization use case with the relevant section of the knowledgebase:

```
1 Optimization:
2   minimize:
3     minimum:
4       algorithms:
5         Random Forest:
6           parameter:
7             config_path: ./configurations/config.yml
8             config_section: General
9             n_estimators:
10            type: int
11            default: 3
12            min: 1
13            max: 100
14          metadata:
15            Class: Surrogate
16            Image: koarch/random_forest:latest
17            Performance: -1
18            Computational Effort: -1
19            RAM usage: -1
20            Min training data: 5
21            input: preprocessed data
```

Code Listing 4 - Knowledge representation of Random Forest

Line 1-3 describe the use case with the desired feature and goal, to optimize the production it is required to minimize the production time. Line 4 starts the algorithm section, which contains several other algorithms. Line 5 and following contain the knowledge representation of the Random Forest algorithm, with parameters, metadata and the required input for the algorithm. Line 6-13 describe the parameters for the algorithm, e.g., where the algorithm finds the configuration files and which sections are relevant but also which algorithm specific parameters can be set and the corresponding range and default values. Line 14-23 contain the metadata for the algorithm. The image specifies the container image for algorithm instantiation. Line 17-22 describe the resource usage of the algorithm and the quality of the results as well as the efficiency of the algorithm. Those values will be adjusted as soon as use case specific results are available. Line 23 specifies how many data points are usually required for training before the algorithm is able to make reliable predictions. Line 24 indicates the necessary

³⁹ Jan Strohschein Github Repository - Knowledge Module:
https://github.com/janstrohschein/KOARCH/tree/master/Big_Data_Platform/Docker/Knowledge
retrieved 10-11-2020

input for the Random Forest algorithm. Another section in the knowledge base contains the information about all preprocessing algorithms, that are able to provide the required input data. Those preprocessing algorithms also define the expected input data. Following the requirements towards the data back to the data sources enables the dynamic creation of data processing pipelines.

There are several possibilities to interact with the knowledge API, i.e., retrieving or updating information for the current use case, access or replace the complete knowledge base, filtering the knowledge base and return the information for a specific use case or algorithm.

All routes and functions defined in `./src/knowledge.py`` are accessible through other modules or via the web interface:

- (GET) Returns the use case information: `/use_case/`
- (PUT) Updates the use case information: `/use_case/`
- (GET) Returns the complete knowledge base: `/knowledge/`
- (PUT) Imports a YAML file as new knowledge base: `/import_knowledge/`
- (GET) Filters the knowledge for a specific use case: `/knowledge/usecase/`
- (GET) Filters the knowledge for a specific algorithm: `/knowledge/algorithm/`
- (GET) Filters the knowledge and returns feasible pipelines: `/knowledge/feasible_pipelines/`
- (GET) Exports the current content of the knowledge base into a YAML file: `/export_knowledge/`

6.4.7 Cognition

The Cognition is an integral part of CAAI that enables the system to dynamically instantiate additional machine learning pipelines, conduct experiments during the production process and evaluate the results to find the most promising algorithms and best production parameters for the current environment. It accesses the information from several data sources and topics to achieve this functionality, as shown in Figure 51 below:

- Knowledge API provides the initial information about available algorithms and the use case and will receive additional performance metrics for the algorithms during production
- Model application topic contains messages with the optimization results from each trained model for the given use case
- Monitoring topic carries messages with real-time information about the production process in order to compare the predicted results with the actual results
- The Cognition sends the new production parameters to the Adaption topic, where the new parameters are applied to the CPPS controller

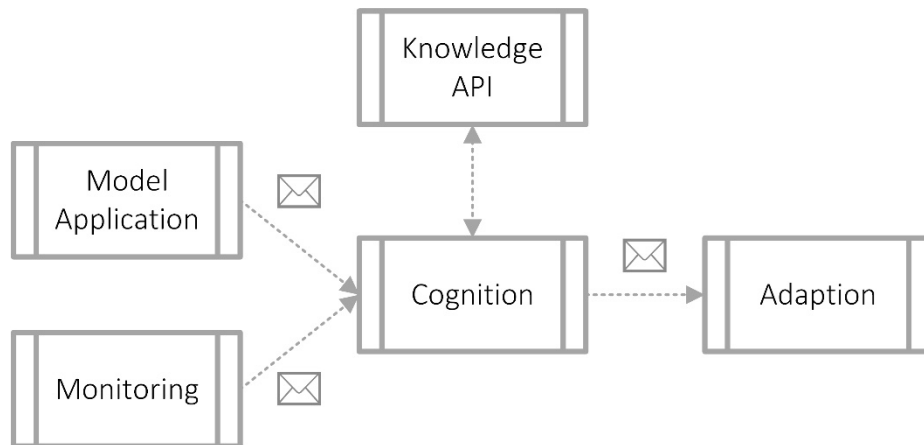


Figure 51 - Cognition communication overview

The complete workflow to process the production data and select an algorithm to adjust the production parameters can be described as follows:

1. During startup the Cognition creates an initial design with parameter values that will be tried in production and sends those values to the Adaption so that the values get applied to the CPPS.
2. The Monitoring module observes the production process and sends the resulting production data to the Cognition.
3. Several models are trained in parallel with the incoming production data. Models can require different minimal amount of training samples to produce promising results. The models are transferred to the Model Application after training.
4. The Model Application tries to find the optimum for the given model through an optimizer and sends the results to the Cognition.
5. The Cognition implements several strategies to find the best solution, ranging from a simple strategy that selects the model with the best predicted outcome to a more complex strategy that calculates the model quality and resource consumption. Thus, it selects the most promising solution and sends it towards the Adaption.
6. The Adaption transfers the new parameter settings to the CPS and the next production cycle starts.
7. The Monitoring module observes the results from the changed parameter settings and reports the new production data to the Cognition.
8. The Cognition compares the predicted results with the real results and calculates the difference. The model quality will be adjusted accordingly and now the real results will be taken into consideration for the next selection process over the predicted results.

The software building block is shown below in Figure 52 and a working example can be found in the Github Repository ⁴⁰. The Cognition will also be described in more detail in the upcoming implementation sections in 6.5.2 and 6.5.3.

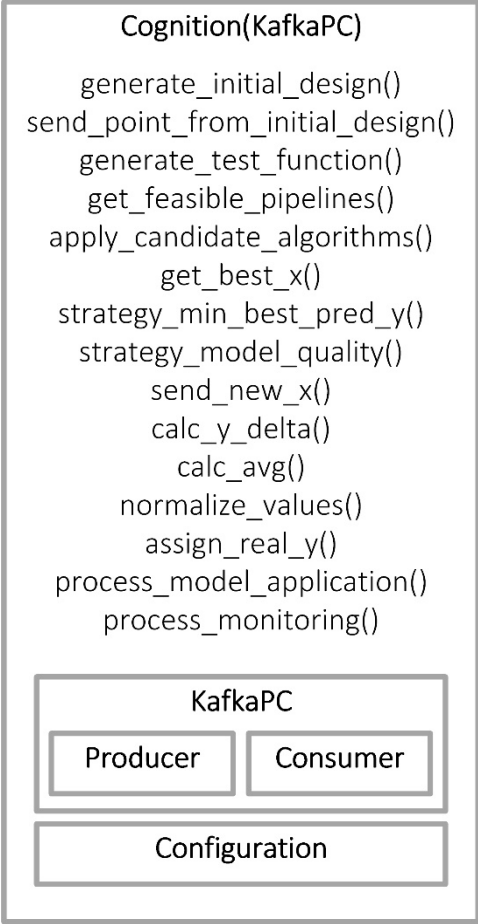


Figure 52 - Cognition Building Block

6.5 USE CASE IMPLEMENTATION

The CAAI architecture was developed and refined in several iterations. Each of the presented use cases has additional requirements and more technical depth.

The first iteration verifies the foundation of the BDP. A single machine learning pipeline collects the status updates of Twitter users, processes the data and extracts the emotions with a pre-defined neural network. All the components of the pipeline run as virtualized Docker containers and communicate with each other via Kafka messages. The virtualization combined with the messaging

⁴⁰ Jan Strohschein Github Repository - VPS Use Case:
https://github.com/janstrohschein/KOARCH/blob/master/Use_Cases/VPS_Popcorn_Production/Docker/readme.md retrieved 22.11.2020

approach enables easy parallelization of heavy workloads and is used for the classification of the status updates.

The second iteration evaluates several machine learning pipelines in parallel to optimize the popcorn production in a VPS. The BDP generates models with Random Forest and Kriging algorithm and uses an optimizer to improve the solution. A first implementation of the Cognition module receives the results as well as the resource usage and selects the winning algorithm, which is allowed to adjust the parameters for the next production cycle. The process data from the next iteration is used to update both models and make the next prediction. Even though the Cognition operates on a fixed set of algorithms, it decides which algorithm should be used to adapt the VPS.

The Cognition module gets more control over the algorithms that are used to optimize the popcorn production in the third iteration. The big data platform is implemented in Kubernetes, which orchestrates the available Docker containers. The introduction of a Knowledge module provides the cognition with information about algorithms that are feasible for a given use case. The Cognition is now able to select candidate algorithms and instantiate them via Kubernetes. Therefore the Cognition can create its own experiments and learn from the results.

The following sections present each of the use cases in more detail.

6.5.1 Detecting Emotions in Social Media

The use case implements an automated approach to study the emotions of a larger group of social media users over time using CAAI. The use case implementation with instructions how to collect and classify tweets is available on Github⁴¹. Further details on the use case can be found in the corresponding paper (Strohschein et al., 2019)

It is possible to extract emotions of social media users from the text of their status updates as shown by Colneric and Demsar, and Tasoulis et al. (Colneric & Demsar, 2018; Tasoulis et al., 2018). This analysis is based on the work of Colneric and Demsar, who were kind enough to publish the resulting machine learning model on GitHub⁴². They utilized deep learning of neural networks to generate a model that is able to detect emotions in English language. Neural networks are a supervised machine learning method and therefore the data needs annotations for the algorithm to learn from. As the authors learned on a massive dataset of 73 billion tweets it was infeasible to manually annotate the dataset. The authors exploited hashtags as annotations, an approach that was successfully used in several other natural language processing studies for sentiment classification (Go et al., 2009; Kouloumpis et al., 2011; Nodarakis et al., 2016), detecting sarcasm (Bamman & Smith, 2015; González-Ibáñez et al., 2011), studying personality traits (Plank & Hovy, 2015) and classifying emotions (Mohammad & Kiritchenko, 2015). As hashtags are selected by the author of a tweet they work well as indicators of their emotions. The machine learning algorithm analyzes the given text and tries to derive the hashtag as target variable.

Emotions can be modeled in a multitude of ways and popular emotion classification schemes were created by Paul Ekman, Robert Plutchik and Douglas McNair along with Maurice Lorr and Leo Droppleman (Ekman, 1999; McNair et al., 1971; Plutchik, 1982). The classification for this analysis is

⁴¹ Jan Strohschein Github Repository - Emotion Analysis Use Case: https://github.com/janstrohschein/KOARCH/tree/master/Use_Cases/other/Social_Media_Emotion_Detection retrieved 24.11.2020

⁴² Niko Colneric Github Repository: <https://github.com/nikicc/twitter-emotion-recognition> retrieved 24.11.2020

done with Ekman's scheme of basic emotions as it covers fear, disgust and anger. The six basic emotions are explained by Ekman and Cordaro (Ekman & Cordaro, 2011) in a later article as follows:

- **Anger:** the response to interference with our pursuit of a goal we care about. Anger can also be triggered by someone attempting to harm us (physically or psychologically) or someone we care about. In addition to removing the obstacle or stopping the harm, anger often involves the wish to hurt the target.
- **Fear:** the response to the threat of harm, physical or psychological. Fear activates impulses to freeze or flee. Often fear triggers anger.
- **Surprise:** the response to a sudden unexpected event. It is the briefest emotion.
- **Sadness:** the response to the loss of an object or person to which you are very attached. The prototypical experience is the death of a loved child, parent, or spouse. In sadness there is resignation, but it can turn into anguish in which there is agitation and protest over the loss and then return to sadness again.
- **Disgust:** repulsion by the sight, smell, or taste of something; disgust may also be provoked by people whose actions are revolting or by ideas that are offensive.
- **Happiness:** feelings that are enjoyed, that are sought by the person. There are a number of quite different enjoyable emotions, each triggered by a different event, involving a different signal and likely behavior. The evidence is not as strong for all of these as it is for the emotions listed above.

The concrete architecture for this use case is shown in Figure 53 below. The BDP connects to Twitter, displayed at the very bottom, to retrieve users and their status updates and sends the results to the Data Bus. Modules in the DPL operate on the incoming data and publish the results onto the Analytics Bus. The user can interact with the big data platform via a Jupyter⁴³ notebook web interface, shown at the top, that provides an interactive analysis.

⁴³ Jupyter is an open-source project that “supports interactive data science and scientific computing across all programming languages”. For more information see: <https://jupyter.org/> retrieved 04.12.2020

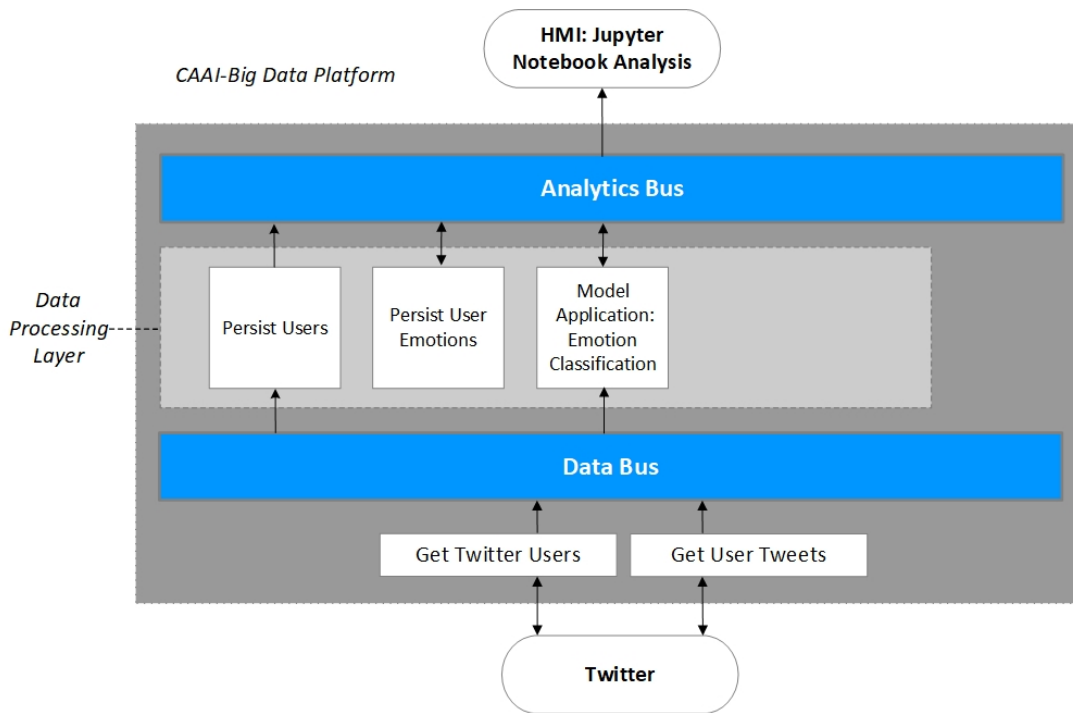


Figure 53 - Concrete CAAI Architecture for Emotion Detection Use Case

The architecture allowed the implementation of a highly modular real-time machine learning pipeline, depicted in Figure 54. The combination of standard building blocks, as presented in 6.4, and new modules for use case specific functionality, e.g., a machine learning model that extracts emotion from status updates, reduced the overall implementation effort. Using an event-based system that communicates through messages allowed to process incoming data in real-time and distribute work-intensive processing onto multiple worker instances.

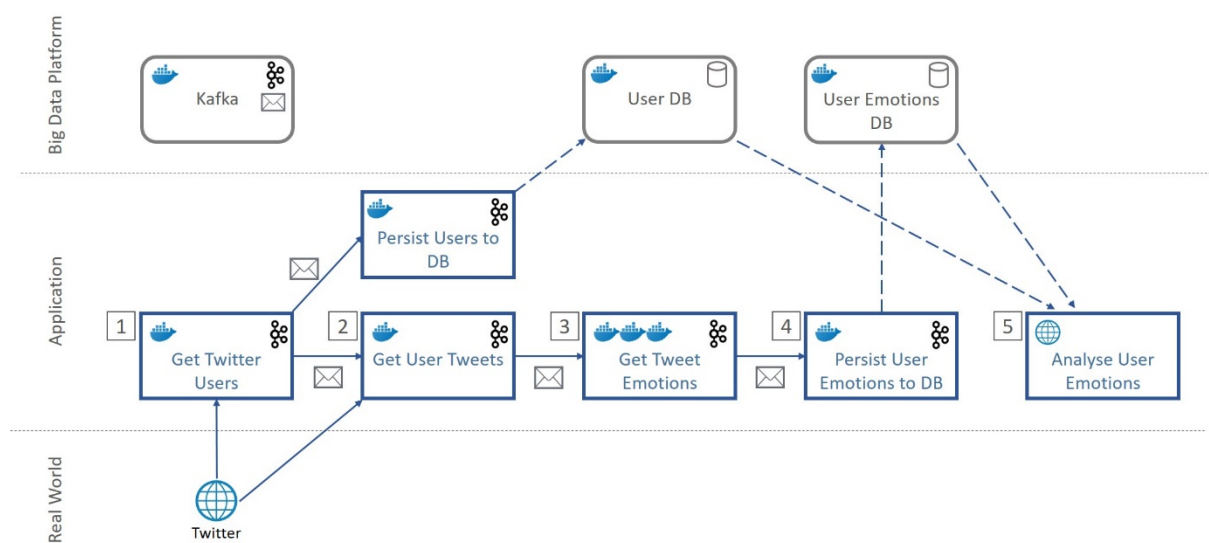


Figure 54 - Data collection and analysis pipeline

In the first step trending Twitter topics are observed and active users, who write status updates on these topics, extracted. A first test-run found that commercial users or algorithms also watch and target these trending topics with their advertising status updates, therefore a high percentage of collected users were duplicates and it was necessary to implement a blacklist in order to analyze every user just once. For every user also their following users were extracted to increase the user sample-size. In a second step for every user the tweets were collected through the Twitter API. So called “retweets”, where a user quotes the status update of someone else, were not collected to analyze only tweets that the user wrote himself. The API imposes limits on the amount of requests per time interval and the number of status updates that can be retrieved for any given user. Therefore, only the last ~3.200 status updates can be collected for each user. Surprisingly this turned out to be an important constraint, because for very active users this was just a fraction of their status update history. The data collection process (step 1 + 2) was able to retrieve over 10 million status updates from 5000 users in a time-span of two weeks. The pre-trained model from Colneric and Demsar was used in the third step to analyze the status updates regarding expressed emotions. The result is a multi-class classification with an allocated percentage value for each emotion, as shown in Table 4 below. This step was run in parallel as the process was very time-consuming and the results were reconciled and persisted in a database (step 4) for further analysis (step 5). The analysis evaluates the emotions of all users over time. Consecutively the users have been grouped for further analysis based on their number of status updates and their followers.

To perform the analysis the two collected and related data sets, i.e., the users and their tweets with the associated emotional classification, are joined to increase the available information. The two tables below describe the datasets and the available fields. Table 8 contains the metadata regarding the Twitter users and the two attributes that are used in the analysis to create user groups, i.e., the number of status updates a user has published and the amount of followers a certain user has.

| User Feature | Description |
|--------------------------------|---|
| user_id (integer) | Automatically generated identification number for every user. |
| user_name (string) | The user can choose the name to display. |
| user_location (string) | The user can specify his/her location. |
| account_created_at (timestamp) | The system records the day and time of account creation. |
| statuses_count (integer) | The amount of status updates a user has written, including retweets. |
| favorites_count (integer) | The number of Tweets this user has liked. |
| followers_count (integer) | The number of followers this account currently has. |
| friends_count (integer) | The number of users this account is following. |
| verified (boolean) | If the user’s identity has been verified by twitter the value is “True”, otherwise “False”. |

Table 8 - Twitter User Features and Descriptions

Table 9 consists of the metadata for a given status update, e.g. which user published the Tweet at which time, but also the text itself and the classification results for each emotion.

| Tweet Analysis Feature | Description |
|-------------------------------|--|
| status_id (integer) | Automatically generated identification number for every tweet. |
| user_id (integer) | Automatically generated identification number for every user. |
| status_created_at (timestamp) | The system records the day and time of tweet creation. |
| text (string) | The tweet text written by the user and analyzed for emotions. |
| anger (float) | The calculated percentage value for a particular emotion in a users' tweet based on the basic emotions model by Ekman. |
| disgust (float) | |
| fear (float) | |
| joy (float) | |
| sadness (float) | |
| surprise (float) | |

Table 9 - Tweet Analysis Features and Descriptions

An example classification that contains only the relevant columns is shown below in Table 10. The status update regards a sport event where an athlete plays his first game for the team and the detected prevalent emotion is joy, followed by surprise.

| Text | Anger | Disgust | Fear | Joy | Sadness | Surprise |
|--|--------------|----------------|-------------|------------|----------------|-----------------|
| 'Anthony Davis makes his debut with the Hornets dropping 21 points and grabbing 7 rebounds.' | 0.0109 | 0.0028 | 0.0425 | 0.7640 | 0.0551 | 0.1244 |

Table 10 - Tweet Analysis multi-class Classification Example

The emotion classification has subsequently been performed for each of the 10 million status updates to investigate the development of emotions expressed by users on the social media platform. The collected user data and the emotion classifications allow to create various aggregations. Table 11 displays the average emotions of all users per year with joy and surprise as the dominant emotions. It is noticeable that sadness is declining while anger and fear are rising.

| Year | Anger (%) | Disgust (%) | Fear (%) | Joy (%) | Sadness (%) | Surprise (%) |
|-------------|------------------|--------------------|-----------------|----------------|--------------------|---------------------|
| 2015 | 3,7 | 1,9 | 14,7 | 32,1 | 14,6 | 33,0 |
| 2016 | 4,5 | 1,9 | 16,4 | 32,8 | 12,6 | 31,8 |
| 2017 | 5,1 | 1,6 | 16,5 | 34,9 | 11,7 | 30,2 |
| 2018 | 5,4 | 1,8 | 17,6 | 35,1 | 11,4 | 28,7 |
| 2019 | 5,8 | 2,0 | 18,0 | 32,9 | 11,8 | 29,6 |

Table 11 - Average Emotions of all Users per Year

It is also possible to create different user groups based on other attributes, e.g., the amount of status updates or the number of followers. Table 12 shows the average emotions for the users with the least amount of followers, while Table 13 displays the average emotions for the users with the highest amount of followers. The distribution of emotions with user groups based on the amount of followers the users have, shows clear distinctions between the groups. Users with the least amount of followers (<25%) express their emotions strongly, they have the highest mean and median values for the emotions anger, fear and joy. Values for those emotions decrease for each of the following user groups. The emotions disgust, sadness and surprise are expressed more strongly the more followers a user has, with the highest values for the user group with more than 75% of followers.

| Year | Anger (%) | Disgust (%) | Fear (%) | Joy (%) | Sadness (%) | Surprise (%) |
|------|-----------|-------------|----------|---------|-------------|--------------|
| 2015 | 6,6 | 0,9 | 19,6 | 36,7 | 8,6 | 27,6 |
| 2016 | 6,7 | 0,9 | 18,6 | 34,3 | 7,5 | 32,0 |
| 2017 | 6,2 | 1,0 | 18,6 | 37,4 | 8,5 | 28,3 |
| 2018 | 6,4 | 1,0 | 19,4 | 37,5 | 7,9 | 27,9 |
| 2019 | 7,9 | 1,0 | 20,9 | 35,9 | 7,6 | 26,7 |

Table 12 – Emotions of users with less than 25% of followers

| Year | Anger (%) | Disgust (%) | Fear (%) | Joy (%) | Sadness (%) | Surprise (%) |
|------|-----------|-------------|----------|---------|-------------|--------------|
| 2015 | 3,5 | 2,1 | 14,6 | 31,5 | 15,3 | 32,9 |
| 2016 | 3,9 | 2,1 | 15,5 | 32,2 | 13,9 | 32,4 |
| 2017 | 4,4 | 2,1 | 15,3 | 32,7 | 13,9 | 31,6 |
| 2018 | 5,1 | 2,1 | 16,4 | 33,6 | 13,3 | 29,5 |
| 2019 | 5,5 | 2,3 | 17,1 | 30,6 | 13,7 | 31,2 |

Table 13 – Emotions of users with more than 75% of followers

The implementation for this use case demonstrates the feasibility of the architecture's technical foundations, i.e., modular development of independent application containers that communicate asynchronously via messages and create a machine learning pipeline. The container pipeline operated under continuous high load without failure to collect and analyze 10 million tweets over the time span of two weeks.

6.5.2 Popcorn production on Versatile Production System

The second use case considers the VPS that has been introduced during the requirement collection in 5.2.1. Several machine learning pipelines work in parallel for online optimization of the popcorn production. A first implementation of the cognitive module evaluates the algorithm results and selects the algorithm that is allowed to adjust the parameters for the next production cycle. The detailed description of the use case and all results can be found in the International Journal of Advanced

Manufacturing Technology (Fischbach et al., 2020). A running demonstration of the implementation can be found on Github⁴⁴.

The VPS is a modular production system, which processes corn to produce popcorn that is used as packaging material. Typically, there are four VPS units, namely delivery, storage, dosing, and production. Due to its modularity, the first three units can be exchanged or removed easily. Depending on the current orders, different configurations are used. The need for different configurations rises, e.g., if a small and exact amount of popcorn should be produced, which is performed by the dosing unit. However, if larger amounts are requested, it is more efficient to renounce the dosing unit because it is slow, and it generates operation costs. Efficiently operating the VPS is a challenge because many parameters influence the result, e.g., the moisture of the corn, the rate of corn that does not pop, or the amount of corn within the reactor. Since not all parameters can be measured inline, data-driven optimization is a promising method to increase efficiency. Therefore, the CAAI architecture perfectly matches the requirements of the VPS use case.

In this use case, all VPS units are used, and small boxes of popcorn are produced. In each batch, one box of popcorn has to be filled. The overage of popcorn produced in one batch, or not fully filled boxes cannot be used, so it is waste. Optimizing the amount of corn in the reactor, as provided by the dosing unit, is the goal. The optimum is a trade-off between three minimization functions: the energy consumption (f1), the processing time (f2), and the amount of corn needed for a small box (f3). These functions are conflicting to some degree. The result of the optimization is a parameter value for the dosing unit. The parameter x controls the run-time of the conveyor and, therefore, indirectly influences the amount of corn processed. As the given optimization problem can be regarded as relatively simple, we will apply a single objective optimization algorithm and compute a weighted sum of the objectives. This results in the following optimization problem:

$$\min \sum_{i=1}^3 w_i f_i(x); w. r. t w_i > 0 \text{ and } \sum_{i=1}^3 w_i = 1 \quad (1)$$

The scalar weights of the corresponding objectives (w_i) are chosen based on user's preferences. As a default, equal weights are used. The minimum of (1) is a Pareto- optimal solution (Marler & Arora, 2010). The model construction requires sampled data for a set of n values of x , which should ideally depict a representative set of all possible settings, i.e., in a space-filling manner. In this case, the set is generated by evaluating an equidistantly spaced design in the complete parameter range of x . The Cognition will evaluate different surrogate models: Random Forest (Breiman, 2001) and Kriging (Krige, 1952). Kriging is especially suitable for modeling continuous data with few variables and comes with an uncertainty measurement. At the same time, Random Forest is also able to model discrete parameters and computes very fast. Recent examples of applications of Kriging and Random Forest in CPPS scenarios can be found in (Jung et al., 2017; Xing et al., 2018).

Data from the real-world VPS was acquired to evaluate the modeling and optimization. This data consists of 36 production cycles with 12 different settings for the runtime of the conveyor. Based on

⁴⁴ Jan Strohschein Github Account, retrieved 23.07.2020, https://github.com/janstrohschein/KOARCH/tree/master/Use_Cases/VPS_Popcorn_Production/Docker

this data, we trained a model that reflects the real behavior of the VPS and utilize it for further experiments. The three different objectives, i.e., the energy consumption, the processing time, and the amount of corn needed were aggregated by taking the sum of the single objectives multiplied with equal weights of 1/3.

The concrete CAAI implementation for the use case is shown in Figure 55 below. The CPPS, displayed at the very bottom, transfers production data to the Data Bus. The modules in the DPL with a connection to the Data Bus use this data as input, i.e., Kriging and Random Forest are used to train models, which are subsequently published onto the Analytics Bus, while the Monitoring module delivers the raw production data to the Analytics Bus. The Model Application retrieves the model and searches for an optimum for each model, which is transferred back to the Analytics Bus.

Modules in the Conceptual Layer determine the system behavior and allow the user to interact with the BDP through various HMIs. The plotting module graphically presents data collected by the reporting module, i.e., production data, model metrics and the associated resource consumption. Another HMI component allows the user to save the data to a CSV or JSON file and change the algorithm parameter settings in the knowledge module. The Cognition evaluates the results and sends the adjusted parameter values to the Analytics Bus. The Adaption receives the new values and transfers those back to the CPPS Controller.

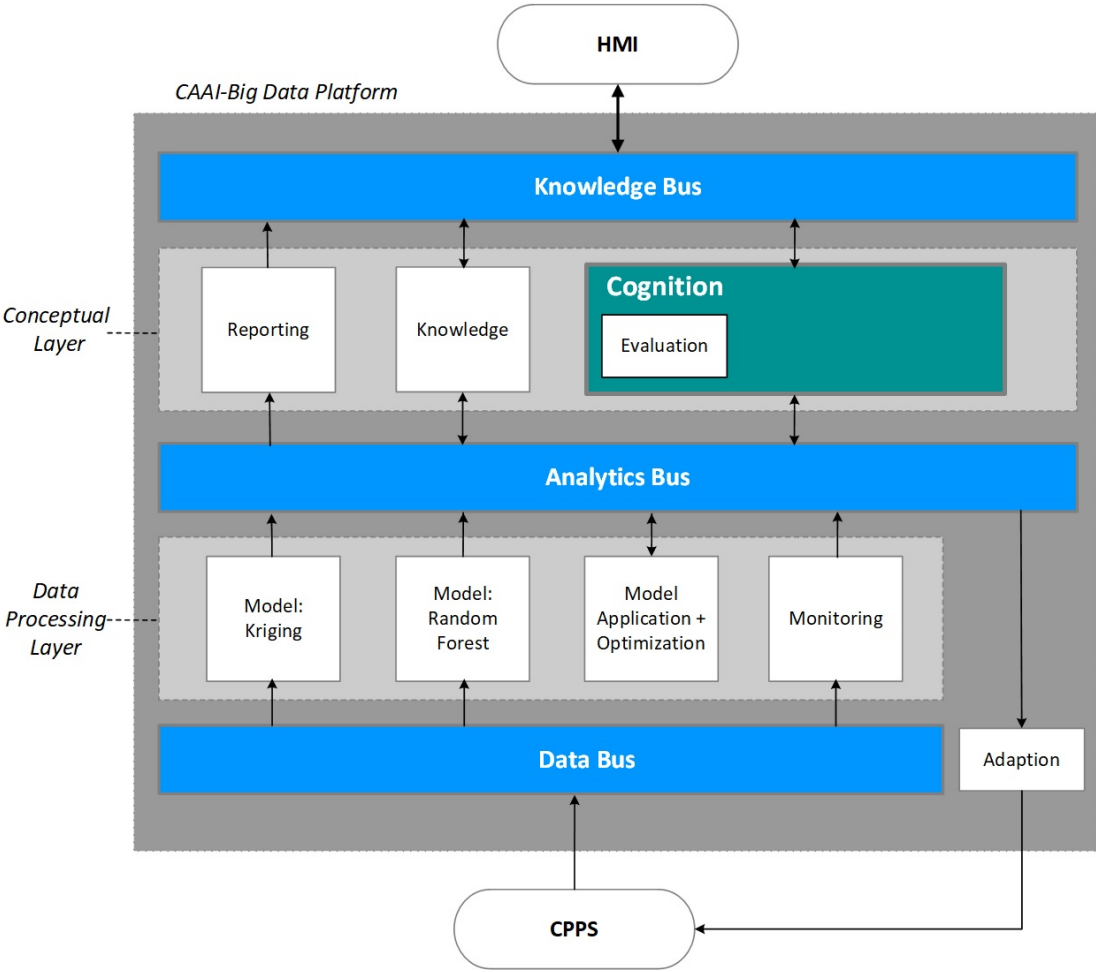


Figure 55 - Concrete CAAI Architecture for Popcorn Production Use Case

The resulting machine learning pipeline is shown in Figure 56 below and considers two phases, i.e., start-up and continuous operation. The system start-up consists of the following steps:

- CAAI creates several containers, one for each module of the CAAI machine learning pipeline.
- The CPPS Module (0) generates a model-based simulation of the production process with data from experiments on the real-world CPPS.
- The Cognition creates the initial design. The initial design consists of 5 points, equally distributed over the search space for x. The Cognition publishes those as starting points to the Analytics Bus, where the Adaption (4) listens.
- The Adaption sends the new parameters to the CPPS Controller (0), where they are retrieved from the CPPS (0) for continuous operation.

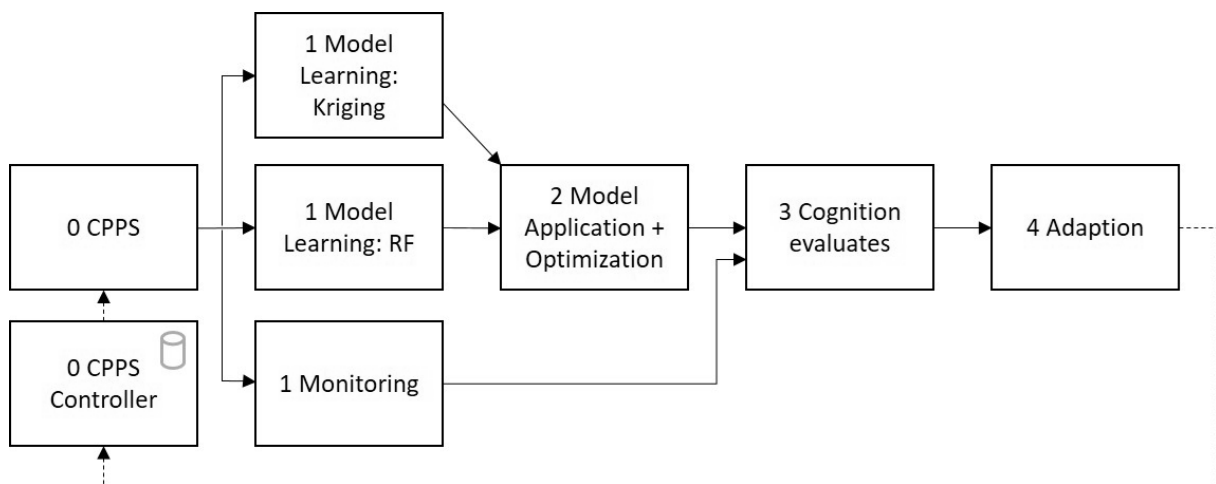


Figure 56 - Popcorn Production Machine Learning Pipeline

When the start-up concludes the continuous operation begins and performs the following operations:

- The CPPS Module (0) uses the simulation to evaluate the incoming points for x and derive the corresponding y-value. The CPPS (0) sends both points to the Data Bus.
- CAAI uses message-based communication so any number of modules can subscribe to a topic and each consumer will receive the data. In this use case the Monitoring (1) and two Model Learning (1) modules listen to new data from the CPPS (0). The Monitoring (1) module transfers the CPPS data from the Data Bus to the Analytics Bus, so the Cognition (3) can evaluate the process data. The Model Learning (1) modules, implementing Kriging and Random Forest algorithms, use leave-one-out cross-validation to calculate RMSE, MAE and R^2 and send those metrics as well as the trained model to the Analytics Bus.
- The Model Application + Optimization (2) module receives both models and uses differential evolution to search for an improved solution. The best predicted y and the corresponding x for each algorithm is published to the Analytics Bus.
- The Cognition (3) is able to evaluate the suggested new values. Once the Cognition decided a new value to try in production, it is send to the Adaption (4), which instructs the CPPS controller and concludes the iteration.

The results have been collected for further analysis of the machine learning pipelines. The comparison of resource consumption, the prediction error and the reached objective function value shows the advantage of the online optimization of the production process with several algorithms working in parallel.

The following plots show the resulting aggregated median values after ten repetitions, each with 20 production cycles. The results do not contain the first five production cycles where the parameter values were determined through the initial design.

Figure 57 plots the CPU consumption in seconds against the VPS production cycles, which is equal to the number of data points used for model training. For both methods, Kriging and Random Forest, an increasing trend can be observed. However, the computation time of Kriging shows a larger slope, compared with Random Forest. Both algorithms behave as expected, stemming from the internal data representation and processing.




Figure 57 - CPU Usage per Production Cycle

The prediction error is the absolute difference between the model prediction of the best found point and the true objective function value. In Figure 58 Kriging shows a nearly constant accurate performance, while Random Forest shows a larger variance and starts to get comparably accurate predictions in the last production cycles.




Figure 58 - Prediction Error per Production Cycle

The reached values of the objective function are depicted in Figure 59. It shows that in the beginning, Kriging outperforms random forest, while later after about 12 cycles, random forests perform comparably to Kriging.




Figure 59 - Objective Function Value per Production Cycle

The use case leads to several interesting conclusions. Online optimization enables the real-time adjustment of the production parameters. The use case implementation with CAAI allows to efficiently train several models in parallel on the production data. However, the models possess very different characteristics. Kriging produces good results with very limited training data, but the computation resources increase quickly. Thus, Kriging is the better choice for the first production cycles. The new production data then allows to improve both models and Random Forest starts to generate good results after roughly 15 production cycles. After 20 production cycles the prediction error and the

objective function value is comparable for both algorithms, while the resource consumption for Random Forest is much lower. Thus, the cognitive module will automatically switch the algorithm that adjusts the production parameters and utilizing more than one algorithm allows taking advantage of each algorithm's strengths.

6.5.3 VPS with cognition on Kubernetes

The third use case extends the previously presented use case that optimizes the popcorn production in a VPS. The second use case demonstrated the online optimization through several algorithms that are trained on the production data in parallel and evaluated by the Cognition to choose the most promising candidate. However, the candidate algorithms were fixed on system start-up as the pure Docker implementation does not allow the dynamic instantiation of additional algorithms. Thus, the third use case implements the BDP on Kubernetes, a technology presented in 6.3.2, to enable dynamic instantiation of additional algorithms through container orchestration. Therefore, the Cognition can create its own experiments and learn from the results.

A detailed description of the use case and all results can be found on Arxiv (Strohschein, Fischbach, et al., 2020). A running example of the implementation is available in the Github repository⁴⁵.

The concrete architecture for the second and the third use case is identical. Thus, Figure 55 in 6.5.2 also shows the concrete CAAI architecture for this use case. However, the BDP is implemented for this use case in Kubernetes, which orchestrates the different modules. Kubernetes enables declarative cluster management, e.g., the user submits a specification of the desired state of an application to Kubernetes, and its controller takes the required measures to reach this state. The user composes the application of different building blocks, which fits the modular approach of the CAAI architecture very well. The smallest building block in Kubernetes is called a pod and consists of one or more application containers and optionally a volume for data storage. The BDP uses two higher-level objects for most of the services, namely the deployment and the job. The deployment is used for all modules of the BDP that require continuous operation, e.g., the cognitive module, the messaging solution, or also the local container registry. A deployment, as shown in Figure 60, specifies the number of replicas of a pod, which the Kubernetes controller instantiates and monitors on the available nodes in the cluster. The controller will start new instances if a single pod or a complete node fails to reach the desired deployment state for the Kubernetes cluster.



Figure 60 - Kubernetes Deployment

The job is meant for one-off execution of a task, e.g., a part of a data processing pipeline. A job, as seen in Figure 61, defines a pod and the desired amount of parallelism or the number of allowed retries, if

⁴⁵ Jan Strohschein Github Repository, retrieved 01.12.2020, https://github.com/janstrohschein/KOARCH/tree/master/Use_Cases/VPS_Popcorn_Production/Kubernetes

the job fails during execution. The Kubernetes controller will track the job progress and manage the whole job lifecycle to free up resources after completion.



Figure 61 - Kubernetes Job

The exemplary Random forest job definition in Code Listing 5 shows the number of retries in line 6, the maximum allowed time for processing in line 7, and the time before a job will be deleted by the controller in line 8. The cognitive module can specify and submit new jobs to dynamically build data processing pipelines. While a job template contains the container image and the command to execute on container start (line 13 - 14), the Cognition can adjust the arguments to pass into the container (line 15), e.g., how many trees the random forest algorithm will create or which is the criterion for a split.

```

1  apiVersion: batch/v1
2  kind: Job
3  metadata:
4    name: Random_Forest
5  spec:
6    backoffLimit: 5
7    activeDeadlineSeconds: 20
8    ttlSecondsAfterFinished: 60
9    template:
10   spec:
11     containers:
12     - name: random_forest
13       image: caai/random_forest
14       command: ["python", "-u", "random_forest.py"]
15       args: ["NumberOfTrees=5"]
16     restartPolicy: OnFailure
  
```

Code Listing 5 – Kubernetes Job Definition

An overview of the complete process to dynamically create a data processing pipeline is depicted in Figure 62. The cognitive component decides which algorithms should be tested on the current use case based on information about available cluster resources from the monitoring module and the knowledge on available algorithms and their properties. The cognition then declares which jobs need to run to form one or more data processing pipelines. The controller subsequently pulls the container images for the given jobs from the container registry and instantiates them.

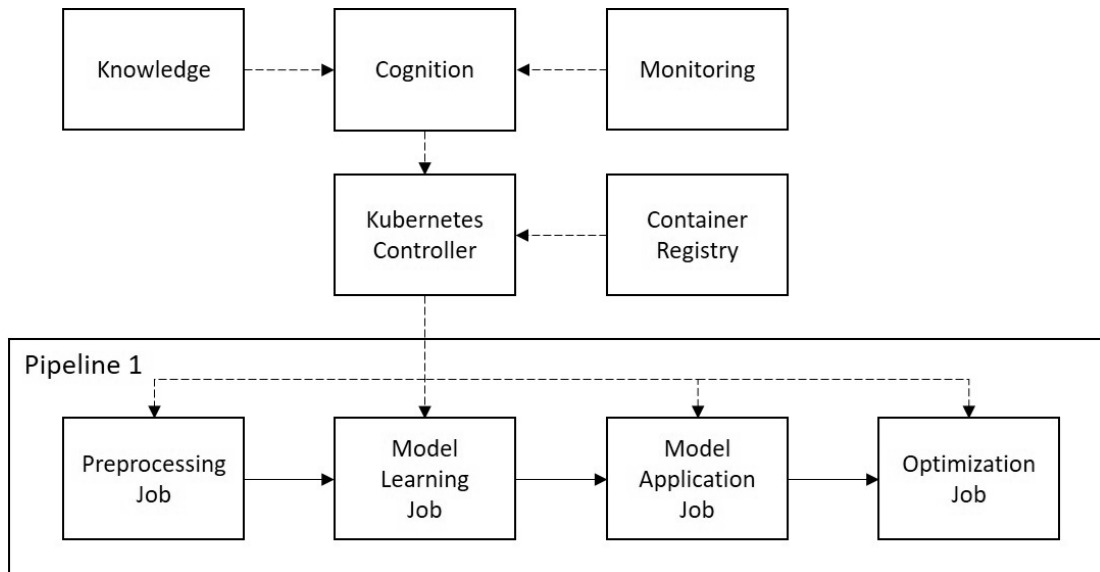


Figure 62 - Cognition dynamically creates new pipelines

However, the decision-making process is not trivial and the algorithm selection through the Cognition is explained in more detail below. The Cognition retrieves information from the Knowledge module about the use case at hand and the available algorithms. The user can specify declarative goals for the use case in a four-step process via the Knowledge API either before system start-up or during production. In the first step, the user selects the overall goal, such as optimization, anomaly detection, condition monitoring, or predictive maintenance. One or more signals that are relevant for the use case are selected in the second step. The third step determines the aggregation functions, e.g., mean, delta, minimum or maximum value. In the last step, the user selects the optimization goal, i.e., minimizing or maximizing. With this four-step process, the user can define the optimization goal for the given use case on an abstract level, as shown in Figure 63.

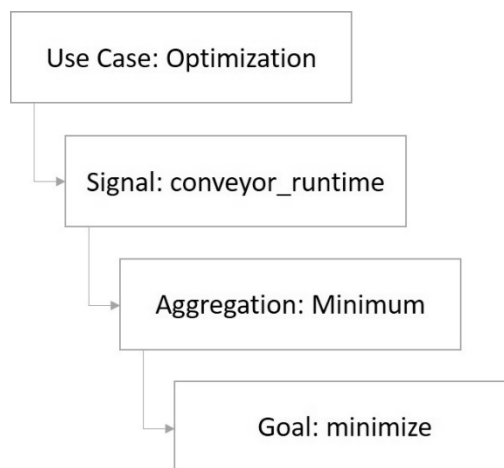


Figure 63 - Declarative goal definition

The Cognition applies the use case knowledge to generate all feasible pipelines and search for promising candidate algorithms, as shown in Code Listing 6 below.

At first, required variables are defined (lines 1-4). Afterwards, the Cognition checks if historical data is available (line 5). If no historical data is available, an initial design representing a list of different parameter configurations is created, the parameters are applied to the CPPS and the resulting data is stored in list d (lines 6-9). Due to the modular design of the architecture, the design method can be easily adapted for different use cases. Currently, a full factorial design is the default configuration. Atkinson et al., Husslage et al., and Morris and Mitchel give an overview of several design methods and corresponding optimality criteria (Atkinson et al., 2014; Husslage et al., 2011; Morris & Mitchell, 1995). If historical production data is available, it will be used as the starting point (line 11).

```

Input: Initial design size  $s$ , selection step size  $\theta$ ,
       algorithm characteristics  $KB$ , goal  $g$ , available
       resources  $r$ , historical data  $dH$ 

1  define List for evaluation  $e$ 
2  define List data  $d$ 
3  define List pipeline resource consumptions  $pr$ 
4  define Parameter  $x$ 
5  if ( $|dH| = 0$ ) then
6  | List parameter  $l \leftarrow$  createInitialDesign( $s$ ,  $KB$ )
7  | forall (parameter  $l$ ) do
8  |    $d \leftarrow d +$  applyToCPPS( $l$ )
9  |    $x = l$ 
10 else
11 |  $d, x \leftarrow dH$ 
12 repeat
13 | if ( $nrIterations \% \theta = 0 \vee \zeta = 1$ ) then
14 |    $\zeta = 0$ 
15 |   testFunctionSet  $S \leftarrow$  generateTestFunctions( $d$ )
16 |   List  $p \leftarrow$  determineFeasiblePipelines( $KB$ ,  $g$ ,  $d$ )
17 |    $p \leftarrow$  selectCandidatePipelines( $KB$ ,  $r$ ,  $e$ )
18 |   forall  $p[i]$  do in parallel
19 |      $e \leftarrow$  applyPipeline( $p[i]$ ,  $S$ )
20 |    $KB, pbest \leftarrow$  ratePipelines( $KB$ ,  $p$ ,  $e$ )
21 |    $xbest \leftarrow$  getBestX( $pbest$ ,  $d$ ,  $x$ )
22 |   if ( $|x - xbest| \geq \epsilon$ ) then
23 |      $e \leftarrow$  applyToCPPS( $xbest$ )
24 |      $x = xbest$ 
25 |    $d \leftarrow d +$  receiveNewData()
26 |   if (  $since \theta = 2$  steps stagnation  $\vee$  performance decrease) then
27 |      $\zeta = 1$ 
28 until true;

```

Code Listing 6 – Cognition Algorithm Selection

After the initialization phase, the algorithm selection process based on data-driven simulation and benchmarking is started within an infinite loop (line 12). The user-defined step size ϑ and a variable ζ ensure that pipelines are changed only after each θ steps, or in situations of prolonged stagnation or performance decrease (line 13 and lines 26-27). In the case that new pipelines shall be selected a test

function set s is generated based on the data d using Gaussian process simulation as presented by Zaefferer and Rehbach (line 15) (Zaefferer & Rehbach, 2020). Simulation intends to reproduce the covariance structure of a set of samples, i.e., the process data, which maintains the topology of the problem landscape. The goal is to analyze the behavior of the performance of candidate algorithms on problems similar to the CPPS problem. The generation of several different instances, either via Spectral simulation or simulation by decomposition (see Zaefferer and Rehbach for details), allows a benchmarking of potentially feasible algorithms. This benchmark analyzes the algorithms regarding their resource consumption and performance, based on their relative rank, not the achieved values. The analysis assumes that the resource consumption depends on the current machine load, the number of function evaluations to perform, and the dimensionality of the problem, but not on the problem landscape structure. Consequently, the resource consumption can be analyzed quite accurately using simulations.

The algorithm description from the knowledge base is used to determine a list of all feasible pipelines in line 16. Based on the available resources, the content of the knowledge base, and possibly existing earlier evaluations, the most suitable pipelines are selected in line 17. The algorithm can exclude pipelines that are known not to be able to calculate a result in the given time. These pipelines are applied in parallel to the test function set s , which is used for parameter tuning and benchmarking of the pipelines (lines 18-19). Instances are drawn randomly for each step, tuning the algorithms first with an equal budget, and subsequently benchmarking on a different instance. The evaluation list e is used in line 20 to rate the pipelines and find the most promising candidate algorithm. As the overall performance measure, a weighted normalized aggregation of the achieved performance on the simulation instance, the memory consumption, and the used CPU time is computed and assigned to each pipeline. The results are used to update the algorithm knowledge in the knowledge base KB. Bartz-Beielstein et al. recently published an overview of best-practices for benchmarking and performance assessment (Bartz-Beielstein et al., 2020). A normalized processing time, computed by dividing the used CPU time by the runtime of a standard algorithm, is suggested by Johnson and McGeoch and Weise et al. (Johnson & McGeoch, 2007; Weise et al., 2014). Thus, a baseline comparator, i.e., random search, is used as a reference algorithm. The goal is, to reward efficient usage of resources, leading to higher accuracy compared to the baseline and relative to the competing algorithms. Simple algorithms with low resource requirements might be good fallback choices when the system load is high and few resources are available, at least if they perform better than the baseline. Consequently, algorithms will be removed from this iteration if the performance is worse than the baseline.

The best pipeline in the remaining list is chosen for application on the CPPS (line 21). If the list is empty, the current x will be returned. Should the new x_{best} differ significantly from the current x , the new x_{best} is applied to the CPPS for the next iteration (line 22-24). The new data produced by the CPPS is appended to the data set d (line 25). If there is no significant improvement after half of the step size, the ζ is set to 1, to perform an unscheduled algorithm selection cycle (lines 26-27). This intends to help in the first production cycles to recover from poor decisions.

The cognitive module was evaluated on production data from the real-world CPPS. A data set was used that repeats 12 different settings for the conveyor belt three times each, for a total of 36 production cycles. This data was used to fit a Gaussian process model as ground-truth. A conditional simulation was used to retrieve an accurate simulation of the popcorn production, so consequently, the model respects the data points. Whereas the test instances, which will be used to perform the benchmark experiments, will use unconditional simulation.

A diverse set of optimization algorithms was selected based on work by Stork et al. and evaluated on the ground-truth and the simulations in R version 3.6.3 (Stork et al., 2018):

- Gaussian process simulations from the COBBS package (version 1.0.0)
- Differential evolution from the DEoptim package (version 2.2-5)
- Generalized simulated annealing from the GenSA package (version 1.1.7)
- Kriging and Random Forest from the SPOT package (version 2.0.6)
- L-BFGS-B from the stats package included in R

The algorithm parameters, the associated ranges, and default values are summarized in Table 14. Some parameter values were chosen according to the relatively low budget of function evaluations, which is suitable for the given use case. This is especially the case for parameters that, e.g., control the number of candidate solutions per iteration.

| Algorithm Family | Algorithm | Parameter | Range | Default |
|------------------|-------------------------|------------|-----------------|---------|
| Population | Differential evolution | popsize | N+ | 5 |
| | | strategy | {1, 2, 3, 4, 5} | 2 |
| | | F | [0, 2] | 0.8 |
| | | CR | [0, 1] | 0.5 |
| | | c | [0, 1] | 0.5 |
| Trajectory | Generalized SA | temp | N+ | 100 |
| | | qv | R | 2.5 |
| | | qa | R | -1 |
| Surrogate | Kriging (SPOT) | designSize | N+ | 7 |
| | | designType | {Lhd, Uniform} | Lhd |
| Surrogate | Random Forest (SPOT) | nrTrees | N+ | 500 |
| | | designSize | N+ | 7 |
| | | designType | {Lhd, Uniform} | Lhd |
| Hill-climber | L-BFGS-B | Imm | N+ | 5 |
| Baseline | Uniform Random Sampling | | | |

Table 14 - Chosen Optimizers with parameter ranges and default values

A correlation between the algorithm performance ranks on the ground-truth and the simulations could be proven experimentally with Pearson’s correlation analysis (Stigler, 1989). The correlation was both high and significant at a correlation coefficient of 0.823 and a low p-value (2.2e-16).

Figure 64 depicts the increase in the mean of the CPU usage of the algorithms on the ground-truth objective over ten repetitions each, including the objective function evaluation. The CPU consumption of the random search is used as the baseline. The development of the memory consumption is not shown as it behaved in the same manner, i.e., the surrogate-based optimizers showed a linear increase in memory consumption, and the other algorithms use a negligible amount of memory. The evaluation showed that the best-performing algorithms also consumed significantly higher resources. Optimizers with low resource demands achieved significantly worse results, even though these algorithms may become feasible when other algorithms are no longer able to compute results within a single production cycle. Consequently, optimizers performing worse than the baseline will be excluded from the selection process, as the baseline itself is a cheap fallback.




Figure 64 - CPU consumption of different optimization algorithms on the ground-truth

A multi-objective performance assessment of algorithms is highly encouraged since resource consumption and solution quality mutually influence each other (Bossek et al., 2020). Thus, CPU and memory consumption are integrated into the overall performance rating. The overall performance of the algorithms on the simulations will be evaluated as follows. The random search performance according to the achieved objective function value, the consumed memory, and the processing CPU time is used as a reference for the competitors. Achieved objective function values are computed as a relative improvement compared to the baseline value. The algorithm will be removed if there is no improvement. The memory and CPU consumption for all remaining algorithms will be divided by the baseline reference values. Each of the factors will be normalized. Thus, the best performing optimizer gets the value 1, and the worst gets the value 0. The other optimizers are scaled in between those based on the individual results. Multiplying every normalized factor with the factor-weight calculates an aggregated performance value. The result for a high weight on the objective function (0.8) and small weights for both resource measures (0.1) is shown in Figure 65.




Figure 65 - Aggregated performance ranks for different optimizers on the simulations

The surrogate-based optimizers Kriging and Random forest consistently achieve the highest ranks (higher is better). Random forest is more resource-efficient and thus ranked higher than Kriging, when both achieve similar objective function values. In scenarios with limited resources, e.g., with an objective function weight of 0.5 and a performance measure weight of 0.25 each, Differential evolution or GenSA achieved the same rank as Kriging, even if they achieve significantly worse objective function values. The result is not meant to prove the general superiority of any algorithm. It is a very use-case and system-specific ranking of the algorithms and suggests which algorithm might perform well in that specific production setting.

The evaluation demonstrates how CAAI uses the cognitive module to automatically select and tune algorithms to perform online optimization for the popcorn production in the VPS. Depending on the difficulty and complexity of the problem at hand, the data-driven simulation of the real-world problem allows an efficient benchmark of the available algorithms and enables the selection of the most suitable one. The implementation of CAAI on Kubernetes provides several advantages. The modular architecture allows the easy modification of the algorithm selection strategy. The user can extend the pool of available algorithms during production via the Knowledge API, and Kubernetes allows to dynamically instantiate the selected algorithms in machine learning pipelines during the production process.

6.6 CAAI EVALUATION

The implementation of various use cases in the previous section demonstrates the generalizability of the CAAI big data platform for the implementation of AI in the I4.0 context. The implementation of the big data platform with Docker and Kubernetes ensures continuous operation and creates a fault-tolerant system. The modular building blocks of the architecture can be used in different use cases and provide a high degree of re-usability. The Cognition chooses from a collection of algorithms and generates feasible pipelines for the given use case. However, the Cognition is able to switch algorithms during operation, if another algorithm becomes more suitable.

In this section the CAAI architecture is systematically evaluated with the same requirements as the existing architectures in 5.4 based on the implemented use cases in 6.5. Table 15 presents the overall evaluation results. The rows represent the requirements, whereas the columns display the results for the five reviewed architectures. A “-” indicates that the requirement is not addressed by the architecture, while “O” means that the architecture partially fulfills a requirement, e.g., there may be additional effort required. Finally, if the architecture fulfills a requirement, the “✓” is used.

| Requirement | Description | CAAI |
|-------------|--|------|
| R.1 | Receive declarative user goals | ✓ |
| R.2 | Specified interfaces are well defined | ✓ |
| R.3 | Strategies to select a suitable algorithm | ✓ |
| R.4 | The system learns from experiences | ✓ |
| R.5 | Knowledge representation including expert knowledge | O/✓ |
| R.6 | Acquire data from heterogeneous and distributed systems via versatile protocols | O/✓ |
| R.7 | Store and manage process data and models | ✓ |
| R.8 | Perform preprocessing | O/✓ |
| R.9 | Learn a model from data with limited time and resources | ✓ |
| R.10 | Perform model analysis or simulation with a limited response time | ✓ |
| R.11 | Interaction with the user, i.e., human-machine-interface | ✓ |
| R.12 | Decision making derives actions, e.g., decide to send new control parameters to the CPPS | ✓ |
| R.13 | Apply the action on the controller | O/✓ |

Table 15 - Evaluation of the requirements for CAAI

The assessment of individual requirements is explained in more detail below:

- R.1: Declarative Goals can be defined by the user before startup of the CAAI machine learning system or during runtime via the associated API. The user defines the goals with a four step grammar that specifies the use case, the relevant signals, the feature and the related feature goal. Rating: ✓
- R.2: CAAI relies on well-defined interfaces as it is a highly modular architecture. All architecture modules communicate via messages. Each message has to be serialized with the associated schema or the message will not be accepted by the broker. Thus, CAAI encourages and enforces the implementation of well-defined interfaces. Rating: ✓

- R.3: The CAAI architecture currently offers two different strategies for algorithm selection, i.e., selecting the best predicted solution and a more complex strategy that considers the predicted solution and the related resource consumption, weighs them according to the user preferences and compares the quality of the predicted solution with the real outcome after the next production cycle. The modular approach allows the user to easily implement additional strategies for algorithm selection. Rating: ✓
- R.4: The system is able to learn from experiences as the environment changes through the online learning approach, e.g., when the material properties change or a tool deteriorates. Online learning is necessary to create a cognitive, reactive system and a distinct advantage over traditional machine learning or even more recent AutoML⁴⁶ methods that process historic production data in batches, isolated from the actual production process. Rating: ✓
- R.5: The CAAI architecture currently includes knowledge about the already available algorithms, but it is also possible to expand the knowledge base. Possible extensions are the properties of additionally implemented algorithms or use case specific knowledge, such as the production machinery specification and parameter constraints. Adding that knowledge presents a certain effort for each use case and requires the knowledge of a domain expert. However, it improves the results of the CAAI machine learning pipelines. Rating: O/✓
- R.6: CAAI supports data acquisition from various sources, such as local files, the internet, Kafka messages from external brokers and OPC UA⁴⁷, a standardized communication protocol for industrial machinery. The modular architecture allows to connect further data sources, but those connectors have to be implemented once. Rating: O/✓
- R.7: CAAI stores process data and machine learning models on the three presented bus systems. Additional management of the data is possible via the Postgres DB and the FastAPI building blocks, that implement common CRUD-operations and allow the user to download the data as JSON or CSV file. Rating: ✓
- R.8: Preprocessing data in real-time can be implemented through the CAAI architecture. The provided building blocks allow to easily connect new modules to the message-driven architecture. However, preprocessing is not automatic and highly use-case-specific. Thus, it benefits tremendously from an expert's domain knowledge. Rating: O/✓
- R.9: The application of several machine learning algorithms in parallel requires learning models under time and resource constraints. Thus, the CAAI architecture implements an algorithm selection strategy that considers the necessary resources, i.e., processing time and memory. Rating: ✓
- R.10: Model evaluation is done with time constraints as results are most valuable if the production parameters can be adjusted within the timeframe of a production cycle. Rating: ✓
- R.11: The CAAI architecture currently provides basic ways to interact with the user. The reporting module aggregates and delivers the process data via a plotting component and FastAPI. Thus, it is possible to retrieve data programmatically or via web interface as JSON or CSV file. The plotting component presents the data via a dynamic website, that shows the time series data in real-time, as soon as new messages arrive. Further enhancements can be more refined dashboards that are more customized for the given use case. Rating: ✓

⁴⁶ AutoML stands for “automated machine learning”, more details on the university website: <https://www.automl.org/> retrieved 16.11.2020

⁴⁷ OPC UA is a platform independent framework to represent and access an information model, more details on the official website: <https://opcfoundation.org/about/opc-technologies/opc-ua/> retrieved 16.11.2020

- R.12: The Cognition of the CAAI architecture evaluates the results of the various machine learning pipelines and will instruct the Adaption module to send new control parameters to the CPPS, if one of the pipelines could improve the current solution. Rating: ✓
- R.13: The current version of the CAAI architecture allows to send control parameters to the CPPS via OPC UA, a standardized communication protocol for industrial applications. The modular architecture allows to instruct the CPPS via other communication protocols. However, these must be implemented manually once. Rating: O/✓

The evaluation shows that CAAI addresses all the requirements from 5.2.5 even though not all requirements are marked as fulfilled yet. However, it has to be noted that the architecture enables the user to make these use case specific extensions and adapt the system to the individual circumstances.

6.7 CONCLUSION

This chapter introduces the CAAI architecture, both the general structure and the technologies used to implement it. Afterwards the created standard building blocks and their application in different use cases are presented. Finally the CAAI architecture is evaluated with the same catalog of requirements that was initially used to assess the existing reference architectures for I4.0.

It turns out that the CAAI architecture meets most requirements, although in some cases additional implementation effort is necessary, e.g., for the integration of domain knowledge or use case specific data pre-processing.

The implementation of the big data platform with Docker and Kubernetes ensures continuous operation and creates a fault-tolerant system. Furthermore, the architecture provides a set of standard building blocks that can be used in different use cases and provide a high degree of re-usability. The successful implementation of different use cases with the described technologies and software building blocks demonstrates the generalizability of the big data platform. However, the modular architecture and standardized interfaces support the extension with additional algorithms and other data processing modules. Overall, extending a set of existing building blocks to meet the use case specific demands leads to a reduced implementation effort in contrast to a solution that is created from the very beginning.

Therefore it can be concluded that the CAAI architecture provides a reliable and extensible system for the application of AI algorithms in the I4.0 context.

7 FUTURE RESEARCH

The publication of the CAAI architecture as open-source software and the implementation in various use cases shows the progress of the big data platform and usefulness of the architecture and the developed standard building blocks. However, there are still many ways the general architecture and the software implementation can be improved.

The decision making process of the Cognitive module can still be refined. The Cognition groups algorithms into various classes based on the provided algorithm knowledge and creating an extensive algorithm topology is currently a hot research topic for related fields such as AutoML. Implementing different strategies to evaluate the machine learning pipelines or efficiently tune promising candidate pipelines is another topic for future research. However, as the CAAI is a modular architecture, those parts of the architecture can easily be replaced with updated modules.

The KOARCH research project that developed the CAAI architecture includes several industry partners that will implement the architecture in additional use cases. This will further prove the generalizability of the architecture as the implementation effort between different use cases should be minimized. Implementing the architecture in more use cases will also provide more insights into possible areas of improvements.

The “Innovation Hub Oberbergischer Kreis” is a regional association of small and medium-sized enterprises that was first introduced in chapter 4.3 as the members participating in the I4.0 questionnaire. They currently evaluate the CAAI architecture for usage in their projects and a further expansion of the cooperation is intended. The Cologne Bonn Airport also implemented the architecture successfully in a pilot project to better predict the landing time of incoming airplanes, and a continuation project is planned for 2021. Wider adoption of the CAAI architecture will further increase its usefulness as it is an open-source project, and users can share their created modules with the community, which decreases future implementation efforts for all users.

The possible integration or collaboration with other proposed reference architectures, such as the IIRA and RAMI4.0, would be another interesting research topic. Both architectures represent a bigger picture as mechanical and electrical engineering viewpoints are included. While the level of abstraction is too high to effectively support the software implementation, it may be fruitful to integrate the CAAI architecture into the respective implementation viewpoints.

It would also be very interesting to accompany and support the introduction process of Industry 4.0 technology in a small or medium-sized manufacturing company. The introduction process concerns many different departments, from leadership to production, HR, IT and logistics. Getting insights into the various perspectives on the introduction process will highlight the challenges small and medium-sized enterprises face. However, assisting in the design of new work processes and getting actual worker feedback would be a valuable starting point for further research. Observing the introduction process will also help to derive best practices for the technical implementation but also the ethical introduction process, which are requested by the companies that participated in the “Industry 4.0 Questionnaire for SMEs”, but are currently lacking.

8 CONCLUSION

The thesis at hand motivates the need for a big data reference architecture to apply AI in I4.0. It presents the CAAI architecture as a possible solution, especially for small and medium-sized enterprises.

Initially, the introduction of AI, big data, and I4.0 establish common ground. The intersection of those exponential technologies holds enormous economic potential, and they catalyze each other. AI algorithms achieve the best results when trained on massive datasets. Big data technology enables the collection and processing of vast amounts of data. I4.0 applications profit from AI and big data as the costs for sensors and processing power rapidly decrease. Existing processes can be optimized, and entirely new business models emerge.

Subsequently, the economic and ethical implications of AI, big data, and I4.0 are examined. These technologies show the potential to revolutionize the global economy and our working lives. Therefore, many national research programs and business ventures emerge in a race of economic powers for global competitive advantages, even though the ethical consequences are not yet completely foreseeable. Workers fear the loss of their jobs and the independence of their work. However, employee support is one of the most important criteria for a successful implementation. Therefore an ethical introduction process is presented, which is including and assisting rather than overburdening the employees.

Large corporations develop most I4.0 technologies and standards. Reasons for this are different technical preconditions and tight budget or missing know-how in small and medium-sized enterprises. The “Industry 4.0 Questionnaire for SMEs” was conducted to gain better insight into those companies’ requirements and needs. The results show that SMEs that collaborate with a bigger partner have a higher technology acceptance rate than SMEs who do not collaborate with a big company. The majority of participants also stated that they need assistance to formulate an I4.0 strategy and evaluate the available I4.0 technologies. They are looking for standards or best practices, which unfortunately are not yet available.

Reference architectures are introduced as they represent best practices for designing a software system and support companies during implementation. An evaluation of the proposed reference architectures for I4.0 shows the possibilities but also the existing shortcomings.

Subsequently, the CAAI architecture for the application of AI in CPPS is presented as a counter-proposal with a focus on software implementation for CPPS. The architecture overview explains all modules, technologies, and their interaction. The CAAI architecture and the related software implementation are open-access and open-source, which creates transparency and leads to a democratization of technology. The available set of standard architecture building blocks eases the implementation effort and assists companies in the I4.0 introduction.

Different use cases demonstrate the applicability of the architecture. The first use case validates the technical foundations of the BDP, but solves the use case with a fixed machine learning pipeline. The two following use cases incorporate algorithm knowledge and implement the Cognition module that instantiates and evaluates new machine learning pipelines with increasing flexibility.

The subsequent evaluation uses the catalog of requirements for I4.0 use cases, that was also used to assess the proposed reference architectures, to verify the capabilities and functionality of the CAAI architecture. Most of the examined requirements can be fulfilled through the architecture, while some requirements need additional use case specific implementation effort, e.g., to integrate process knowledge from a domain expert.

The resulting modular architecture can be implemented on existing IT resources or through various cloud computing providers. Thus, a small or medium-sized enterprise can use the architecture to optimize production processes even if their budget does not allow them to buy the required hardware outright. The architecture supports the users through various already implemented building blocks and easy extension with new modules. Thus, the implementation effort is low in comparison to an AI solution for the application in the I4.0 context that is built from scratch.

Wider adoption of the CAAI architecture will further increase its usefulness as it is an open-source project and users can share their created modules with the community which decreases future implementation effort for all users.

APPENDIX A - INDUSTRY 4.0 QUESTIONNAIRE FOR SMES

Industry 4.0 Readiness for SMEs

Bitte wählen Sie eine Sprache aus. / Please choose a language. / Por favor elige un idioma.

- Deutsch
 English
 Español

Welcome to the questionnaire on the introduction of Industry 4.0 in small- and medium-sized enterprises conducted by Jan Strohschein and Heide Faeskorn-Woyke from Technische Hochschule Köln and Ana María Lara-Palma from Universidad de Burgos. Your responses help us to assess the current status of introduction in small- and medium-sized enterprises and identify needs for further research.

Completing the survey will take you only 15-20 minutes. Your responses to this assessment are anonymous and strictly confidential. Your answers will only be reported in the aggregate and individual responses are not shared.

Q1: In which country is your company located? *

Q2: Please estimate the size of your company's domestic workforce. *

- less than 10 employees
 10 to 49 employees
 50 to 249 employees
 more than 250 employees

Q3: Please estimate the number of employees working on the I4.0 introduction in the following departments. *

| | 0 - 5 | 5 - 10 | >10 |
|-------------------------------|-----------------------|-----------------------|-----------------------|
| Leadership / Company Strategy | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Manufacturing | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| IT | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| HR | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Q4: Do you plan to increase the number of employees working on the I4.0 introduction in the following departments. *

| | yes | no |
|-------------------------------|-----------------------|-----------------------|
| Leadership / Company Strategy | <input type="radio"/> | <input type="radio"/> |
| Manufacturing | <input type="radio"/> | <input type="radio"/> |
| IT | <input type="radio"/> | <input type="radio"/> |
| HR | <input type="radio"/> | <input type="radio"/> |

Q5: Please estimate your 2019 revenues. *

- up to 2 million euros
- 2 million up to 10 million euros
- 10 million up to 50 million euros
- Not specified

Industry 4.0 Introduction

Q6-Q13: Do you agree with the following statements regarding the introduction of I4.0 in your company? *

| | strongly disagree | disagree | neither agree nor disagree | agree | strongly agree |
|--|-----------------------|-----------------------|----------------------------|-----------------------|-----------------------|
| "Introducing I4.0 in our company is a good idea." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "Our company is well prepared to introduce I4.0." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "Our company adopts the I4.0 strategy of a (bigger) partner." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "Our company feels the need to introduce I4.0 to stay competitive." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "Our company feels the need to introduce I4.0 to continue collaboration with (bigger) partners." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "For our work the usage of I4.0 is important." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "The benefits of I4.0 introduction are well known to our company and clearly evaluated." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "The costs of I4.0 introduction are well known to our company and clearly evaluated." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Strategy and organization

Industry 4.0 is about more than just improving existing products or processes through the use of digital technologies – it actually offers the opportunity to develop entirely new business models. For this reason, its implementation is of great strategic importance.

Q14: How would you describe the implementation status of your Industry 4.0 strategy? *

- No strategy exists
- Pilot initiatives launched
- Strategy in development
- Strategy formulated
- Strategy in implementation
- Strategy implemented

Q15: Do you use indicators to track the implementation status of your Industry 4.0 strategy? *

- Yes, we have a system of indicators that we consider appropriate
- Yes, we have a system of indicators that gives us some orientation
- No, our approach is not yet that clearly defined

Q16: Which technologies do you already use in your company? *

- Sensor technology
- Mobile end devices
- RFID
- Realtime location systems
- Big data to store and evaluate real-time data
- Cloud technologies as scalable IT infrastructure
- Embedded IT systems
- M2M communications
- None of the above

Q17.1: In which parts of your company have you invested in the implementation of Industry 4.0 in the past two years, and what are your plans for the future?

-- Investments in the past 2 years -- *

| | Large | Medium | Small | None |
|--------------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Research and development | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Production/manufacturing | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Purchasing | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Logistics | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Sales | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Services | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| IT | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Q17.2: In which parts of your company have you invested in the implementation of Industry 4.0 in the past two years, and what are your plans for the future?

-- Investments the next 5 years -- *

| | Large | Medium | Small | None |
|--------------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Research and development | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Production/manufacturing | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Purchasing | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Logistics | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Sales | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Services | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| IT | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Q18: In which areas does your company have systematic technology and innovation management? *

- IT
- Production technology
- Product development
- Services
- Centralized, in integrative management
- No systematic technology and innovation management

Industry 4.0 Technology

Industry 4.0 transforms manufacturing through integration of the digital into the physical world. A multitude of emerging new technologies and concepts require evaluation through a company to assess the individual benefits.

Q19: How well do you know the following Industry 4.0 related concepts? *

| | no knowledge | little knowledge | good knowledge | very good knowledge |
|-----------------------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| Smart Factory | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Digital Twin | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Smart Workpiece | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Horizontal / Vertical Integration | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Artificial Intelligence | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Data-driven Services | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Q20: Do you plan to implement the following Industry 4.0 related concepts? *

| | not planned | planned | implementation ongoing | implementation complete |
|-----------------------------------|-----------------------|-----------------------|------------------------|-------------------------|
| Smart Factory | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Digital Twin | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Smart Workpiece | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Horizontal / Vertical Integration | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Artificial Intelligence | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Data-driven Services | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Q21: Do you wish for easier access to supporting materials and knowledge about the following Industry 4.0 related concepts? *

| | yes | no | undecided |
|-----------------------------------|-----------------------|-----------------------|-----------------------|
| Smart Factory | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Digital Twin | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Smart Workpiece | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Horizontal / Vertical Integration | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Artificial Intelligence | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Data-driven Services | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Employees

Employees help companies realize their digital transformation and are the ones most affected by the changes of the digital workplace. Their direct working environment is altered, requiring them to acquire new skills and qualifications. This makes it more and more critical that companies prepare their employees for these changes through appropriate training and continuing education.

Q22-Q28: Do you agree with the following statements regarding the effects on employees by introduction of I4.0 in your company? *

| | strongly disagree | disagree | undecided | agree | strongly agree |
|--|-----------------------|-----------------------|-----------------------|-----------------------|-----------------------|
| "Using I4.0 would improve our employees job performance." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "Our employees like the idea of introducing an I4.0 system." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "The interaction between employees and I4.0 systems will be clear and understandable." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "Clear interactions between employees and I4.0 systems are important." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "Learning how to use a I4.0 system will be easy for our employees." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| "Our employees face the new I4.0 challenges with confidence." | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Q29: How do you assess the skills of your employees when it comes to the future requirements under Industry 4.0? *

| | Non-existing | Existent, but inadequate | Adequate | Not relevant |
|---|-----------------------|--------------------------|-----------------------|-----------------------|
| IT infrastructure | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Automation technology | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Data analytics | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Data security / communications security | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Development or application of assistance systems | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Collaboration software | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |
| Non-technical skills such as systems thinking and process understanding | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> | <input type="radio"/> |

Q30: Are you making efforts to acquire the skills that employees are currently lacking? Through special training seminars, knowledge transfer systems, coaching, etc. *

- no training planned
- training planned
- training ongoing
- training completed

Survey completed

Thank you very much for your participation.

Please click "Done" to submit the survey.

BIBLIOGRAPHY

- Aaker, D. A. (1991). *Managing brand equity: capitalizing on the value of a brand name*. Free Press.
- Abel, J., Hirsch-Kreinsen, H., Steglich, S., & Wienzek, T. (2019). *Akzeptanz von Industrie 4.0*.
- Adadi, A., & Berrada, M. (2018). Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI). *IEEE Access*, 6, 52138–52160. <https://doi.org/10.1109/ACCESS.2018.2870052>
- Adolphs, P., Bedenbender, H., Dirzus, D., Ehlich, M., Epple, U., Hankel, M., Heidel, R., Hoffmeister, M., Huhle, H., Kärcher, B., Koziol, H., Pichler, R., Pollmeier, S., Schewe, F., Walter, A., Waser, B., & Wollschlaeger, M. (2015). Referenzarchitekturmodell Industrie 4.0 (RAMI4.0). In *VDI /VDE Statusreport* (Vol. 0, Issue April).
- Airaksinen, A., Luomaranta, H., Alajääskö, P., & Roodhuijzen, A. (2015). Statistics on small and medium-sized enterprises. *Eurostat*, September, 1–14. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Statistics_on_small_and_medium-sized_enterprises&oldid=463558
- Anderson, J. R. (1996). A Simple Theory of Complex Cognition. In *American Psychologist* (Vol. 51, Issue 4, pp. 355–365). <https://doi.org/10.1037/0003-066X.51.4.355>
- Angelov, S., Grefen, P., & Greefhorst, D. (2012). A framework for analysis and design of software reference architectures. *Information and Software Technology*, 54(4), 417–431. <https://doi.org/10.1016/j.infsof.2011.11.009>
- Apache Software Foundation. (2020). *Apache Hadoop 3.3.0 - MapReduce Tutorial*. <https://hadoop.apache.org/docs/r3.3.0/hadoop-mapreduce-client/hadoop-mapreduce-client-core/MapReduceTutorial.html>
- Appleton, B. (2000). *Patterns and Software: Essential Concepts and Terminology*. <http://www.bradapp.net/lastmodified02/14/2000>
- Arias-Oliva, M., Pelegrín-Borondo, J., Lara-Palma, A. M., & Juaneda-Ayensa, E. (2020). Emerging cyborg products: An ethical market approach for market segmentation. *Journal of Retailing and Consumer Services*, 55, 102140. <https://doi.org/10.1016/j.jretconser.2020.102140>
- Arsov, I., & Watson, B. (2019). *Potential Growth in Advanced Economies*.
- Atkinson, A. C., Fedorov, V. V., Herzberg, A. M., & Zhang, R. (2014). Elemental information matrices and optimal experimental design for generalized regression models. *Journal of Statistical Planning and Inference*, 144(1), 81–91. <https://doi.org/10.1016/j.jspi.2012.09.012>
- Auschitzky, E., Hammer, M., & Rajagopaul, A. (2014). How big data can improve manufacturing. *McKinsey & Company Inc.*, 2(July), 1–4. <https://www.mckinsey.com/business-functions/operations/our-insights/how-big-data-can-improve-manufacturing#>
- Bach, M. P., Bertonecel, T., Meško, M., Vugec, D. S., & Ivančić, L. (2020). Big data usage in european countries: Cluster analysis approach. *Data*, 5(1), 1–17. <https://doi.org/10.3390/data5010025>
- Bagheri, B., Yang, S., Kao, H. A., & Lee, J. (2015). Cyber-physical systems architecture for self-aware machines in industry 4.0 environment. *IFAC-PapersOnLine*, 28(3), 1622–1627. <https://doi.org/10.1016/j.ifacol.2015.06.318>
- Bamman, D., & Smith, N. A. (2015). Contextualized sarcasm detection on twitter. *Proceedings of the 9th International Conference on Web and Social Media, ICWSM 2015*, 574–577.
- Bangemann, T., Bauer, C., Bedenbender, H., Diesner, M., & Epple, U. (2015). Industrie 4.0 – Technical

- Assets; Grundlegende Begriffe, Konzepte, Lebenszyklen und Verwaltung. *VDI/VDE-Gesellschaft Mess- Und Automatisierungstechnik (GMA)*, November.
- Banko, M., & Brill, E. (2001). *Scaling to very very large corpora for natural language disambiguation*. 26–33. <https://doi.org/10.3115/1073012.1073017>
- Barr, A., & Feigenbaum, E. A. (1981). *The handbook of artificial intelligence. Volume I*. HeurisTech Press.
- Bartz-Beielstein, T., Doerr, C., Bossek, J., Chandrasekaran, S., Eftimov, T., Fischbach, A., Kerschke, P., Lopez-Ibanez, M., Malan, K. M., Moore, J. H., Naujoks, B., Orzechowski, P., Volz, V., Wagner, M., & Weise, T. (2020). Benchmarking in Optimization: Best Practice and Open Issues. *ArXiv*. <http://arxiv.org/abs/2007.03488>
- Bass, L., Clements, P., & Kazman, R. (1998). *Software Architecture in Practice*. Addison Wesley. https://books.google.de/books/about/Software_Architecture_in_Practice.html?id=uKhQAAAAAAAJ&redir_esc=y
- Bauer, K., Diemer, J., Claus, H., Lowen, U., & Michels, S. J. (2017). Benefits of Application Scenario:- Value-Based Service. In *Plattform Industrie 4.0*. https://www.plattform-i40.de/I40/Redaktion/EN/Downloads/Publikation/benefits-application-scenario.pdf?__blob=publicationFile&v=6
- Bauernhansl, T. (2017). Die Vierte Industrielle Revolution – Der Weg in ein wertschaffendes Produktionsparadigma. In *Handbuch Industrie 4.0 Bd.4* (pp. 1–31). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-662-53254-6_1
- Baum, G. (2013). *Innovationen als Basis der nächsten Industrierevolution* (pp. 37–53). https://doi.org/10.1007/978-3-642-36917-9_3
- Bessen, J. (2018). *AI and Jobs: The Role of Demand*. <http://www.nber.org/papers/w24235>
- Best, B. J., & Lebiere, C. (2003). Teamwork, Communication, and Planning in ACT-R Agents Engaging in Urban Combat in Virtual Environments. *Proceedings of the 2003 IJCAI Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions*, 64–72. <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.90.5632&rep=rep1&type=pdf>
- Binfield, K. (2004). *Writings of the Luddites*. <https://jhupbooks.press.jhu.edu/title/writings-luddites>
- Bishop, C. M. (2006). *Pattern Recognition and Machine Learning*.
- Bitkom, VDMA, & ZVEI. (2015). Umsetzungsstrategie Industrie 4.0. - Ergebnisbericht der Plattform Industrie 4.0. *Plattform Industrie 4.0, April*, 100. http://www.plattform-i40.de/sites/default/files/150410_Umsetzungsstrategie_0.pdf
- BMBF. (2018). *Industrie 4.0*. BMBF. <https://www.bmbf.de/de/zukunftsprojekt-industrie-4-0-848.html>
- BMW. (2018). *Relationships between I4.0 Components-Composite Components and Smart Production*. www.bmw.de
- BMW. (2018). *Key points for a Federal Government Strategy on Artificial Intelligence*. 1–9. https://www.bmw.de/Redaktion/EN/Downloads/E/key-points-for-federal-government-strategy-on-artificial-intelligence.pdf?__blob=publicationFile&v=5
- Bonin, H., Gregory, T., & Zieran, U. (2015). Übertragung der Studie von Frey/Osborne (2013) auf Deutschland. In *Zew* (Issue 2013). <https://doi.org/ISSN 0174-4992>
- Bossek, J., Kerschke, P., & Trautmann, H. (2020). A multi-objective perspective on performance

- assessment and automated selection of single-objective optimization algorithms. *Applied Soft Computing Journal*, 88. <https://doi.org/10.1016/j.asoc.2019.105901>
- Bothell, D. (2019). *ACT-R 7.13 Reference Manual*.
- Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/A:1010933404324>
- Bughin, J., Manyika, J., Woetzel, J., Lund, S., Chui, M., Batra, P., Ko, R., & Sanghvi, S. (2017). Jobs lost, jobs gained: Workforce transitions in a time of automation. *McKinsey Global Institute, December*, 1–160. https://www.mckinsey.com/~media/mckinsey/featured_insights/Future_of_Organizations/What_the_future_of_work_will_mean_for_jobs_skills_and_wages/MGI-Jobs-Lost-Jobs-Gained-Report-December-6-2017.ashx
- Buhr, D. (2015). Social Innovation Policy for Industry 4.0. *Friedrich Ebert Stiftung*, 1–24.
- Bullinaria, J. A. (2005). Artificial intelligence: The roots , goals and sub-fields of AI. *Artificial Intelligence and Its Subfield*, 1–16. https://www.google.com/url?sa=t&rct=j&q=&esrc=s&source=web&cd=3&cad=rja&uact=8&ved=2ahUKEwiazOXJ9-TkAhVHQ0EAHd7XB_sQFjACegQIABAC&url=https%3A%2F%2Fwww.cs.bham.ac.uk%2F~jxb%2FIAl%2Fw2.pdf
- Bunte, A., Fischbach, A., Strohschein, J., Bartz-Beielstein, T., Faeskorn-Woyke, H., & Niggemann, O. (2019). Evaluation of Cognitive Architectures for Cyber-Physical Production Systems. *IEEE International Conference on Emerging Technologies and Factory Automation, ETFA, 2019-Septe*, 729–736. <https://doi.org/10.1109/ETFA.2019.8869038>
- Bunte, A., Stein, B., & Niggemann, O. (2019). *Model-Based Diagnosis for Cyber-Physical Production Systems Based on Machine Learning and Residual-Based Diagnosis Models*.
- Burton, W. N., Schultz, A. B., Chen, C.-Y., & Edington, D. W. (2008). The association of worker productivity and mental health: a review of the literature. *International Journal of Workplace Health Management*, 1(2), 78–94. <https://doi.org/10.1108/17538350810893883>
- Bush, V. (1945). As we may think. *The Atlantic Monthly*.
- Campbell, M., Hoane, A. J., & Hsu, F. H. (2002). Deep Blue. *Artificial Intelligence*, 134(1–2), 57–83. [https://doi.org/10.1016/S0004-3702\(01\)00129-1](https://doi.org/10.1016/S0004-3702(01)00129-1)
- Campetelli, A., Irlbeck, M., Bytschkow, D., Cengarle, V., & Schorp, K. (2014). *Reference Framework for the Engineering of Cyber-Physical Systems: A First Approach*.
- Cartledge, C. (2016). *How Many Vs are there in Big Data?* 1–4. <http://clc-ent.com/TBDE/Docs/vs.pdf>
- Chui, M., Manyika, J., Miremadi, M., Henke, N., Chung, R., Nel, P., & Malhotra, S. (2018). Notes from the AI frontier. Insights from hundreds of use cases. *McKinsey Global Institute*, 36. https://www.mckinsey.com/~media/mckinsey/featured_insights/artificial_intelligence/notes_from_the_ai_frontier_applications_and_value_of_deep_learning/notes-from-the-ai-frontier-insights-from-hundreds-of-use-cases-discussion-paper.ashx
- Cloutier, R., Muller, G., Verma, D., Nilchiani, R., Hole, E., & Bone, M. (2010). The Concept of Reference Architectures. *Systems Engineering*, 13(1). <https://doi.org/10.1002/sys>
- Colneric, N., & Demsar, J. (2018). Emotion Recognition on Twitter: Comparative Study and Training a Unison Model. *IEEE Transactions on Affective Computing*, 3045(c). <https://doi.org/10.1109/TAFFC.2018.2807817>

- Confluent. (2020). *Schema Management*. <https://docs.confluent.io/current/schema-registry/index.html>
- Conniff, R. (2014). What the Luddites Really Fought Against. *Smithsonian Magazine*, 1–4. <https://www.smithsonianmag.com/history/what-the-luddites-really-fought-against-264412/>
- Council of Economic Advisers. (2016). *ECONOMIC REPORT OF THE PRESIDENT*. https://obamawhitehouse.archives.gov/sites/default/files/docs/ERP_2016_Book_Complete_JA.pdf
- Crevier, D. (1993). *AI: The Tumultuous History of the Search for Artificial Intelligence*. Basic Books.
- Cummings, T. G., & Manring, S. L. (1977). The relationship between worker alienation and work-related behavior. *Journal of Vocational Behavior*, 10(2), 167–179. [https://doi.org/10.1016/0001-8791\(77\)90053-7](https://doi.org/10.1016/0001-8791(77)90053-7)
- Da Silveira, G., Borenstein, D., & Fogliatto, F. S. (2001). Mass customization: Literature review and research directions. *International Journal of Production Economics*, 72(1), 1–13. [https://doi.org/10.1016/S0925-5273\(00\)00079-7](https://doi.org/10.1016/S0925-5273(00)00079-7)
- Danna, K., & Griffin, R. W. (1999). Health and well-being in the workplace: A review and synthesis of the literature. *Journal of Management*, 25(3), 357–384. <https://doi.org/10.1177/014920639902500305>
- Dauth, W., Findeisen, S., Südekum, J., & Wößner, N. (2017). German Robots – The Impact of Industrial Robots on Workers. In *IAB Discussion Paper*. <http://doku.iab.de/discussionpapers/2017/dp3017.pdf>
- De Laat, P. B. (2018). Algorithmic Decision-Making Based on Machine Learning from Big Data: Can Transparency Restore Accountability? *Philosophy & Technology*, 31, 525–541. <https://doi.org/10.1007/s13347-017-0293-z>
- De Wispelaere, J., & Stirton, L. (2004). The Many Faces of Universal Basic Income. *The Political Quarterly*, 75(3), 266–274. <https://doi.org/10.1111/j.1467-923x.2004.00611.x>
- Dean, J., & Ghemawat, S. (2004). MapReduce: Simplified data processing on large clusters. *OSDI 2004 - 6th Symposium on Operating Systems Design and Implementation*, 137–149. <https://doi.org/10.21276/ijre.2018.5.5.4>
- Decker, A. (2016). Industry 4.0 and SMEs in the Northern Jutland Region. In *Value Creation in International Business: Volume 2: An SME Perspective* (pp. 309–335). Springer International Publishing. https://doi.org/10.1007/978-3-319-39369-8_13
- Deloitte. (2018). The fourth industrial revolution is here: Are you ready? *Deloitte Insight*, 1–26. <https://doi.org/10.1016/j.jbusres.2015.10.029>
- Deloitte Switzerland. (2015). Industry 4.0. Challenges and solutions for the digital transformation and use of exponential technologies. *Deloitte*, 1–30. <https://www2.deloitte.com/content/dam/Deloitte/ch/Documents/manufacturing/ch-en-manufacturing-industry-4-0-24102014.pdf>
- Dhiman, H., & Rocker, C. (2019). Worker Assistance in Smart Production Environments Using Pervasive Technologies. *2019 IEEE International Conference on Pervasive Computing and Communications Workshops, PerCom Workshops 2019*, 95–100. <https://doi.org/10.1109/PERCOMW.2019.8730771>
- DIN. (2000). *DIN EN 61512-1 Batch control*. <https://www.beuth.de/de/norm/din-en-61512-1>

1/24306472

- DIN. (2014). *DIN EN 62264-1 Enterprise-control system integration*. <https://www.beuth.de/de/norm/din-en-62264-1/207270059>
- DIN. (2016). *DIN SPEC 91345 Reference Architecture Model Industrie 4.0 (RAMI4.0)*. <https://doi.org/https://dx.doi.org/10.31030/2436156>
- DIN. (2017). *DIN EN 62890 Life-cycle management for systems and products used in industrial process measurement, control and automation*. <https://www.beuth.de/de/norm-entwurf/din-en-62890/269121145>
- Dirks, K. T., & Ferrin, D. L. (2002). Trust in leadership: Meta-analytic findings and implications for research and practice. *Journal of Applied Psychology*, 87(4), 611–628. <https://doi.org/10.1037/0021-9010.87.4.611>
- Docker. (2020). *Docker Overview*. <https://github.com/docker/docker.github.io/blob/master/get-started/overview.md>
- Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, 55(10), 78. <https://doi.org/10.1145/2347736.2347755>
- Draken, E. (2019). *Comparison of Time-Series Data Transport Formats: Avro, Parquet, CSV*. <https://ericdraken.com/comparison-time-series-data-transport-formats/>
- Drath, R., & Horch, A. (2014). Industrie 4.0: Hit or hype? *IEEE Industrial Electronics Magazine*, 8(2), 56–58. <https://doi.org/10.1109/MIE.2014.2312079>
- Ekman, P. (1999). Basic emotions. In *Handbook of Cognition and Emotion*. John Wiley & Sons Ltd.
- Ekman, P., & Cordaro, D. (2011). What is meant by calling emotions basic. *Emotion Review*, 3(4), 364–370. <https://doi.org/10.1177/1754073911410740>
- Eliot, L. (2020). *Overcoming Racial Bias In AI Systems And Startlingly Even In AI Self-Driving Cars*. <https://www.forbes.com/sites/lanceeliot/2020/01/04/overcoming-racial-bias-in-ai-systems-and-startlingly-even-in-ai-self-driving-cars/#1b73aeea723b>
- Enslow, B. (1989). The payoff from expert systems. *Across the Board*, 54–58. <https://stacks.stanford.edu/file/druid:sb599zp1950/sb599zp1950.pdf>
- Erboz, G. (2018). *HOW TO DEFINE INDUSTRY 4 .0 : The Main Pillars of Industry 4 .0. November 2017*.
- Eskildsen, J. K., & Dahlgaard, J. J. (2000). A causal model for employee satisfaction. *Total Quality Management*, 11(8), 1081–1094. <https://doi.org/10.1080/095441200440340>
- European Commission. (2017a). *Germany: Industrie 4.0*.
- European Commission. (2017b). *Spain: Industria Conectada 4.0*.
- European Commission. (2020a). *Artificial Intelligence | Shaping Europe's digital future*. Ec.Europa.Eu. <https://ec.europa.eu/digital-single-market/en/artificial-intelligence>
- European Commission. (2020b). *National strategies on Artificial Intelligence - A European perspective in 2019 - Country report - Germany*.
- European Commission. (2020c). *National strategies on Artificial Intelligence - A European perspective in 2019 - Country report - Spain*. https://ec.europa.eu/knowledge4policy/ai-watch_en
- European Commission. (2020d). *White Paper on Artificial Intelligence A European approach to*

- excellence and trust*. 27. https://ec.europa.eu/commission/sites/beta-political/files/political-guidelines-next-commission_en.pdf.
- Eurostat. (2018). *Small and medium-sized enterprises: an overview*. <https://ec.europa.eu/eurostat/web/products-eurostat-news/-/EDN-20181119-1>
- Federal Ministry for Economic Affairs and Energy (BMWi). (2019). Which criteria do Industrie 4.0 products need to fulfil? *Platform Industrie 4.0*, 32. www.bmwi.de
- Feldmann, S., Lässig, R., Herweg, O., Rauen, H., & Synek, P.-M. (2017). *Predictive Maintenance - Servicing tomorrow and where we are really at today*. https://industrie40.vdma.org/documents/4214230/17409951/1495448865651_VDMA_Predictive_Maintenance_E.pdf/dcd02e04-3c13-460d-ad27-dcfa2d43c65b
- Ferreira, C. (2015). *Microservices and IBM Bluemix meetup presentation*. <https://www.slideshare.net/fe01134/microservices-and-ibm-bluemix-meetup-presentation>
- Ferreira, F., Faria, J., Azevedo, A., & Marques, A. L. (2016). Product Lifecycle Management Enabled by Industry 4.0 Technology. *Advances in Manufacturing Technology*, 3. <https://doi.org/10.3233/978-1-61499-668-2-349>
- Fischbach, A., Strohschein, J., Bunte, A., Stork, J., Faeskorn-Woyke, H., Moriz, N., & Bartz-Beielstein, T. (2020). CAAI - A Cognitive Architecture to Introduce Artificial Intelligence in Cyber-Physical Production Systems. *The International Journal of Advanced Manufacturing Technology*. <https://doi.org/https://doi.org/10.1007/s00170-020-06094-z>
- Fisher, R. A. (1992). *Statistical Methods for Research Workers* (pp. 66–70). Springer, New York, NY. https://doi.org/10.1007/978-1-4612-4380-9_6
- Frey, C. B., & Osborne, M. (2013). *The future of employment*. 18(1), 19–22. <https://doi.org/10.4274/tjem.2290>
- Frost&Sullivan. (2015). *Industry 4.0 Business Ecosystem — Decoding the New Normal* (Issue September). <https://store.frost.com/industry-4-0-business-ecosystem-decoding-the-new-normal.html>
- Furman, J. (2017). Should We Be Reassured If Automation in the Future Looks Like Automation in the Past? In A. Agrawal, J. Gans, & A. Goldfarb (Eds.), *The Economics of Artificial Intelligence: An Agenda* (pp. 317–328). University of Chicago Press.
- Furman, J., & Robert Seamans. (2018). *AI and Economy*. 33. <http://www.nber.org/papers/w24689>
- Future Image Industry 4.0*. (2012).
- Galster, M., & Avgeriou, P. (2011). *Empirically-grounded Reference Architectures: A Proposal*.
- Gantz, J., & Reinsel, D. (2011). *Extracting Value From Chaos*. <https://fdocuments.in/document/idc-iview-extracting-value-from-chaos-2011-data-storage-etc.html>
- Gerbert, P., Lorenz, M., Rüßmann, M., Waldner, M., Justus, J., Engel, P., & Harnisch, M. (2015). Industry 4.0: The future of productivity and growth in manufacturing industries. *BCG Perspective*, 1–20. https://doi.org/10.1007/978-981-13-3384-2_13
- Ghemawat, S., Gbioff, H., & Leung, S.-T. (2003). *The Google File System*.
- Go, A., Bhayani, R., & Huang, L. (2009). *Twitter Sentiment Classification using Distant Supervision*. 1–6. <https://www-cs.stanford.edu/people/alecmgo/papers/TwitterDistantSupervision09.pdf>

- Gobierno de Espana. (2019). *Spanish RDI Strategy in Artificial Intelligence*. <https://cpage.mpr.gob.es>
- González-Ibáñez, R., Muresan, S., & Wacholder, N. (2011). *Identifying Sarcasm in Twitter: A Closer Look*. <http://www.vidarholen.net/contents/interjections/>
- Gorecky, D., Schmitt, M., Loskyll, M., & Zühlke, D. (2014). Human-machine-interaction in the industry 4.0 era. *Proceedings - 2014 12th IEEE International Conference on Industrial Informatics, INDIN 2014*, 289–294. <https://doi.org/10.1109/INDIN.2014.6945523>
- Graetz, G., & Michaels, G. (2018). Robots at work. In *Review of Economics and Statistics* (Vol. 100, Issue 5, pp. 753–768). https://doi.org/10.1162/rest_a_00754
- Greenhaus, J. H., Bedeian, A. G., & Mossholder, K. W. (1987). Work experiences, job performance, and feelings of personal and family well-being. *Journal of Vocational Behavior*, 31(2), 200–215. [https://doi.org/10.1016/0001-8791\(87\)90057-1](https://doi.org/10.1016/0001-8791(87)90057-1)
- Grodzinsky, F. S., Miller, K. W., & Wolf, M. J. (2010). *Developing artificial agents worthy of trust: "Would you buy a used car from this artificial agent?"* <https://doi.org/10.1007/s10676-010-9255-1>
- Haag, S., & Anderl, R. (2018). Digital twin – Proof of concept. *Manufacturing Letters*, 15(June), 64–66. <https://doi.org/10.1016/j.mfglet.2018.02.006>
- Halevy, A., Norvig, P., & Pereira, F. (2009). The unreasonable effectiveness of data. *IEEE Intelligent Systems*.
- Handley, L. (2019). *Amazon beats Apple and Google to become the world's most valuable brand*. CNBC. <https://www.cnbc.com/2019/06/11/amazon-beats-apple-and-google-to-become-the-worlds-most-valuable-brand.html>
- Hankel, M. (2015). Das Referenzarchitekturmodell. *ZVEI*, 1.0(April), 2. http://www.zvei.org/Downloads/Automation/ZVEI-Faktenblatt-Industrie4_0-RAMI-4_0.pdf
- Harmon, P. (2019). *AI Plays Games*. Forbes. <https://www.forbes.com/sites/cognitiveworld/2019/02/24/ai-plays-games/#243a83c54a49>
- Haupt, M. (2016). *Exponential Technology Defined*. <https://michaelhaupt.com/exponential-technology-defined-374e2db882b0>
- Hays, J., & Efros, A. A. (2007). Scene Completion Using Millions of Photographs. In *Computer Graphics Proceedings, Annual Conference Series*. <http://graphics.cs.cmu.edu/projects/scene-completion/>
- Hecklau, F., Galeitzke, M., Flachs, S., & Kohl, H. (2016). Holistic Approach for Human Resource Management in Industry 4.0. *Procedia CIRP*, 54, 1–6. <https://doi.org/10.1016/j.procir.2016.05.102>
- Heienbrok, K. (1997). *Entwicklung und Inbetriebnahme einer Mehrgrößendosierregelung für ein Betonfertigteilwerk*.
- Hengstler, M., Enkel, E., & Duelli, S. (2016). Applied artificial intelligence and trust-The case of autonomous vehicles and medical assistance devices. *Technological Forecasting and Social Change*, 105, 105–120. <https://doi.org/10.1016/j.techfore.2015.12.014>
- Hermann, M., Pentek, T., & Otto, B. (2016). Design principles for industrie 4.0 scenarios. *Proceedings of the Annual Hawaii International Conference on System Sciences, 2016-March*, 3928–3937. <https://doi.org/10.1109/HICSS.2016.488>
- Hightower, K., Burns, B., & Beda, J. (2019). *Kubernetes up & running*. O'Reilly.

- Hsu, F. (2002). *Behind deep blue : building the computer that defeated the world chess champion*. Princeton University Press.
- Hu, L., Xie, N., Kuang, Z., & Zhao, K. (2012). Review of cyber-physical system architecture. *Proceedings - 2012 15th IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing Workshops, ISORCW 2012*, 25–30. <https://doi.org/10.1109/ISORCW.2012.15>
- Hübner, I. (2015). RAMI 4.0 und die Industrie-4.0 Komponente. *Open Automation*, 24–29. <https://www.smart-production.de/open-automation/news-detailansicht/nsctrl/detail/News/rami-40-und-die-industrie-40-komponente-20151017/>
- Husslage, B. G. M., Rennen, G., van Dam, E. R., & den Hertog, D. (2011). Space-filling Latin hypercube designs for computer experiments. *Optimization and Engineering*, 12(4), 611–630. <https://doi.org/10.1007/s11081-010-9129-8>
- IBM. (2016). *Main features and benefits of message queuing*. IBM Knowledge Center. https://www.ibm.com/support/knowledgecenter/SSFKSJ_9.1.0/com.ibm.mq.pro.doc/q002630_.htm
- IEEE Computer Society. (2000). *IEEE 1471-2000 - IEEE Recommended Practice for Architectural Description for Software-Intensive Systems*. <https://standards.ieee.org/standard/1471-2000.html>
- ISO. (2011). *ISO - ISO/IEC/IEEE 42010:2011 - Systems and software engineering — Architecture description*. <https://www.iso.org/standard/50508.html>
- Jiang, J.-R. (2018). An improved cyber-physical systems architecture for Industry 4.0 smart factories. *Special Issue Article Advances in Mechanical Engineering*, 10(6), 1–15. <https://doi.org/10.1177/1687814018784192>
- Johnson, D. S., & McGeoch, L. A. (2007). *Experimental Analysis of Heuristics for the STSP* (pp. 369–443). Springer, Boston, MA. https://doi.org/10.1007/0-306-48213-4_9
- Jung, C., Zaefferer, M., Bartz-Beielstein, T., & Rudolph, G. (2017). Metamodel-based optimization of hot rolling processes in the metal industry. *International Journal of Advanced Manufacturing Technology*, 90(1–4), 421–435. <https://doi.org/10.1007/s00170-016-9386-6>
- Kagermann, H., Anderl, R., Gausemeier, J., Schuh, G., & Wahlster, W. (2016). Industrie 4.0 in a Global Context - Strategies for Cooperating with International Partners. In *Acatech Study*. <http://search.proquest.com/docview/197176054?accountid=8330%5Cnhttp://library.anu.edu.au:4550/resserv?genre=unknown&issn=00333352&title=Public+Administration+Review&volume=67&issue=3&date=2007-05-01&atitle=Federalism+in+a+Global+Context&spage=590&auiast=R>
- Kagermann, H., Wahlster, W., & Helbig, J. (2013). Recommendations for implementing the strategic initiative INDUSTRIE 4.0 - Final report of the Industrie 4.0 Working Group. *Acatech - National Academy of Science and Engineering*, April, 84.
- Kaplan, A., & Haenlein, M. (2019). Siri, Siri, in my hand: Who's the fairest in the land? On the interpretations, illustrations, and implications of artificial intelligence. *Business Horizons*, 62(1), 15–25. <https://doi.org/10.1016/j.bushor.2018.08.004>
- Karimov, J., Rabl, T., Katsifodimos, A., Samarev, R., Heiskanen, H., & Markl, V. (2018). Benchmarking distributed stream data processing systems. *Proceedings - IEEE 34th International Conference on Data Engineering, ICDE 2018*, 1519–1530. <https://doi.org/10.1109/ICDE.2018.00169>

- Kobayashi, Y., Kaneyoshi, A., Yokota, A., & Kawakami, N. (2008). Effects of a Worker Participatory Program for Improving Work Environments on Job Stressors and Mental Health among Workers: A Controlled Trial. *J Occup Health*, *50*, 455–470.
- Koolen, C., & Van Cranenburgh, A. (2017). *These are not the Stereotypes You are Looking For: Bias and Fairness in Authorial Gender Attribution*. Association for Computational Linguistics.
- Kouloumpis, E., Wilson, T., & Moore, J. (2011). Twitter Sentiment Analysis: The Good the Bad and the OMG! *International AAAI Conference on Web and Social Media*.
- Kraemer-Eis, H., & Passaris, G. (2015). SME Securitization in Europe. *The Journal of Structured Finance*, *20*(4), 97–106. <https://doi.org/10.3905/jsf.2015.20.4.097>
- Kreps, J. (2014a). *I Heart Logs: Event Data, Stream Processing, and Data Integration*. <https://books.google.com/books?hl=es&lr=&id=gdiYBAAAQBAJ&pgis=1>
- Kreps, J. (2014b). *Questioning the Lambda Architecture*. O'Reilly. <https://www.oreilly.com/radar/questioning-the-lambda-architecture/>
- Krige, D. G. (1952). A Statistical Approach to Some Basic Mine Valuation Problems on the Witwatersrand. *Journal of the Chemical, Metallurgical and Mining Society of South Africa*, 201–215. <https://doi.org/10.2307/3006914>
- Kruchten, P. (1995). The 4+1 View Model of Software Architecture. *IEEE Software*, *12*(6), 42–50.
- Laird, J. E., & Mohan, S. (2018). Learning fast and slow: Levels of learning in general autonomous intelligent agents. *32nd AAAI Conference on Artificial Intelligence, AAAI 2018, Ebbinghaus 1885*, 7983–7987.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence*, *33*(1), 1–64. [https://doi.org/10.1016/0004-3702\(87\)90050-6](https://doi.org/10.1016/0004-3702(87)90050-6)
- Landers, J. (1981). *Quantification in History, Topic 4: Hypothesis Testing II-Differing Central Tendency*.
- Laney, D. (2001). *3D Data Management: Controlling Data Volume, Velocity, and Variety*. <https://blogs.gartner.com/doug-laney/files/2012/01/ad949-3D-Data-Management-Controlling-Data-Volume-Velocity-and-Variety.pdf>
- Laney, D., & Beyer, M. (2012). *The Importance of “Big Data”: A Definition*. The Importance of “Big Data”: A Definition. <https://www.gartner.com/en/documents/2057415/the-importance-of-big-data-a-definition>
- Larson, B. N. (2017). *Gender as a Variable in Natural-Language Processing: Ethical Considerations*. Association for Computational Linguistics.
- Lee, J., Bagheri, B., & Kao, H. A. (2015). A Cyber-Physical Systems architecture for Industry 4.0-based manufacturing systems. *Manufacturing Letters*, *3*, 18–23. <https://doi.org/10.1016/j.mfglet.2014.12.001>
- Lee, J., Jin, C., & Bagheri, B. (2017). Cyber physical systems for predictive production systems. *Production Engineering*. <https://doi.org/10.1007/s11740-017-0729-4>
- Lehman, J., Laird, J., & Rosenbloom, P. (2006). A GENTLE INTRODUCTION TO SOAR, AN ARCHITECTURE FOR HUMAN COGNITION: 2006 UPDATE. *Science*, *0413013*, 1–37.
- Li, X., Hess, T. J., & Valacich, J. S. (2008). Why do we trust new technology? A study of initial trust formation with organizational information systems. In *Journal of Strategic Information Systems* (Vol. 17, Issue 1). <https://doi.org/10.1016/j.jsis.2008.01.001>

- Lichtblau, K., Stich, V., Bertenrath, R., Blum, M., Bleider, M., Millack, A., Schmitt, K., Schmitz, E., & Schröter, M. (2015). *Industry 4.0 Readiness*.
- Lin, J. (2011). Technological adaptation, cities, and New Work. *Review of Economics and Statistics*, 93(2), 554–574. https://doi.org/10.1162/REST_a_00079
- Lin, S.-W., Mellor, S., Munz, H., & Barnstedt, E. (2017). *Architecture Alignment and Interoperability An Industrial Internet Consortium and Plattform Industrie 4.0 Joint Whitepaper*. https://www.plattform-i40.de/PI40/Redaktion/EN/Downloads/Publikation/whitepaper-iic-pi40.pdf?__blob=publicationFile&v=3
- Lin, S.-W., Miller, B., Durand, J., Bleakley, G., Chigani, A., Martin, R., Murphy, B., & Crawford, M. (2019). *The Industrial Internet of Things Volume G1: Reference Architecture*. <https://www.iiconsortium.org/pdf/IIRA-v1.9.pdf>
- Maier, A., Niggemann, O., & Eickmeyer, J. (2015). On the learning of timing behavior for anomaly detection in cyber-physical production systems. *CEUR Workshop Proceedings*, 1507(August), 217–224.
- Maier, M. (2013). *Towards a Big Data Reference Architecture*. 1–144.
- Manhart, K. (2013). *Industrie 4.0 könnte schon bald Realität sein*. <https://computerwelt.at/knowhow/industrie-4-0-konnte-schon-bald-realitat-sein/>
- Mann, H. B., & Whitney, D. R. (1947). On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics*, 18(1), 50–60. <https://doi.org/10.1214/AOMS/1177730491>
- Manyika, J., Chui Brown, M., B. J., B., Dobbs, R., Roxburgh, C., & Hung Byers, A. (2011). Big data: The next frontier for innovation, competition and productivity. *McKinsey Global Institute*, June, 156. <https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/big-data-the-next-frontier-for-innovation>
- Manyika, J., Chui, M., & Bughin, J. (2013). Disruptive technologies: Advances that will transform life, business, and the global economy. *McKinsey Global ...*, May, 176. [https://www.mckinsey.com/~media/McKinsey/Business Functions/McKinsey Digital/Our Insights/Disruptive technologies/MGI_Disruptive_technologies_Full_report_May2013.ashx](https://www.mckinsey.com/~media/McKinsey/Business%20Functions/McKinsey%20Digital/Our%20Insights/Disruptive%20technologies/MGI_Disruptive_technologies_Full_report_May2013.ashx)
- Manyika, J., Chui, M., & Joshi, R. (2018). Modeling the global economic impact of AI. *McKinsey*, September. <https://www.mckinsey.com/featured-insights/artificial-intelligence/notes-from-the-ai-frontier-modeling-the-impact-of-ai-on-the-world-economy>
- Marler, R. T., & Arora, J. S. (2010). The weighted sum method for multi-objective optimization: New insights. *Structural and Multidisciplinary Optimization*, 41(6), 853–862. <https://doi.org/10.1007/s00158-009-0460-7>
- Marz, N., & Warren, J. (2015). *Big Data: Principles and best practices of scalable realtime data systems*. Manning Publications. <https://www.manning.com/books/big-data>
- Matt, D., & Rauch, E. (2020). *SME 4.0: The Role of Small- and Medium-Sized Enterprises in the Digital Transformation*. Springer International Publishing. <https://doi.org/10.1007/978-3-030-25425-4>
- Matt, D. T., & Rauch, E. (2013). Design of a network of scalable modular manufacturing systems to support geographically distributed production of mass customized goods. *Procedia CIRP*, 12, 438–443. <https://doi.org/10.1016/j.procir.2013.09.075>
- Matzler, K., Hinterhuber, H. H., Friedrich, S. A., & Stahl, H. K. (2003). Core issues in German strategic

- management research. *Problems and Perspectives in Management*, 1(1), 148–160.
- Matzler, K., & Renzl, B. (2006). The relationship between interpersonal trust, employee satisfaction, and employee loyalty. *Total Quality Management and Business Excellence*, 17(10), 1261–1271. <https://doi.org/10.1080/14783360600753653>
- McCarthy, J., Minsky, M. L., Rochester, N., & Shannon, C. E. (2006). A proposal for the Dartmouth summer research project on artificial intelligence. *AI Magazine*, 27(4), 12–14.
- McDermott, J. (1982). R1: A rule-based configurer of computer systems. *Artificial Intelligence*, 19(1), 39–88. [https://doi.org/10.1016/0004-3702\(82\)90021-2](https://doi.org/10.1016/0004-3702(82)90021-2)
- McKnight, D. H., Cummings, L. L., & Chervany, N. L. (1998). Initial Trust Formation in New Organizational Relationships. *The Academy of Management Review*, 23(3), 473. <https://doi.org/10.2307/259290>
- McNair, D. M., Lorr, M., & Droppleman, L. F. (1971). EITS Manual for the Profile of Mood States. *Educational and Industrial Testing Service*, 3(27), 1984. https://soeg.kb.dk/discovery/fulldisplay?docid=proquest614459791&context=PC&vid=45KBDK_KGL:KGL&lang=da&search_scope=MyInst_and_CI&adaptor=Primo_Central&tab=Everything&query=any,contains,McNair DM, Lorr M, Droppleman L.. Manual for the
- Michaelis, M. (2019). *Preparing for the Exponential Technology Revolution*. <https://docs.microsoft.com/en-us/archive/msdn-magazine/2019/november/exponential-technologies-preparing-for-the-exponential-technology-revolution>
- Microsoft. (2018). 2019 Manufacturing Trends Report. In *Microsoft Dynamic 365*. <http://info.microsoft.com/rs/157-GQE-382/images/EN-US-CNTNT-Report-2019-Manufacturing-Trends.pdf>
- Mohammad, S. M., & Kiritchenko, S. (2015). Using Hashtags to Capture Fine Emotion Categories from Tweets. *Computational Intelligence*, 31(2), 301–326. <https://doi.org/10.1111/coin.12024>
- Monostori, L., Kádár, B., Bauernhansl, T., Kondoh, S., Kumara, S., Reinhart, G., Sauer, O., Schuh, G., Sihn, W., & Ueda, K. (2016). Cyber-physical systems in manufacturing. *CIRP Annals - Manufacturing Technology*, 65(2), 621–641. <https://doi.org/10.1016/j.cirp.2016.06.005>
- Moor, J., Minsky, M., & Shannon, C. (2006). Artificial Intelligence Conference : The Next Fifty Years. *AI Magazine*, 27(4), 87–91.
- Moore, G. E. (1965). Cramming more components onto integrated circuits. *Electronics*, 38(8). <https://newsroom.intel.com/wp-content/uploads/sites/11/2018/05/moores-law-electronics.pdf>
- Morris, M. D., & Mitchell, T. J. (1995). Exploratory designs for computational experiments. *Journal of Statistical Planning and Inference*, 43(3), 381–402. [https://doi.org/10.1016/0378-3758\(94\)00035-T](https://doi.org/10.1016/0378-3758(94)00035-T)
- Mosler, W. (1998). Full Employment and Price Stability. In *Journal of Post Keynesian Economics* (Vol. 20, pp. 167–182). Taylor & Francis, Ltd. <https://doi.org/10.2307/4538575>
- Mullarkey, S., Jackson, P. R., Wall, T. D., Wilson, J. R., & Grey-Taylor, S. M. (1997). The impact of technology characteristics and job control on worker mental health. *Journal of Organizational Behavior*, 18(5), 471–489. [https://doi.org/10.1002/\(SICI\)1099-1379\(199709\)18:5<471::AID-JOB810>3.0.CO;2-V](https://doi.org/10.1002/(SICI)1099-1379(199709)18:5<471::AID-JOB810>3.0.CO;2-V)

- Murphy, B., Lin, S.-W., & Burger, B. (2016). *Applying the Industrial Internet Reference Architecture to a Smart Grid Testbed*.
- Nachar, N. (2008). The Mann-Whitney U: A Test for Assessing Whether Two Independent Samples Come from the Same Distribution. *Tutorials in Quantitative Methods for Psychology*, 4(1), 13–20. <https://doi.org/10.20982/tqmp.04.1.p013>
- Nakagawa, E. Y., Becker, M., & Maldonado, J. C. (2012). A knowledge-based framework for reference architectures. *Proceedings of the ACM Symposium on Applied Computing, March*, 1197–1202. <https://doi.org/10.1145/2245276.2231964>
- Narkhede, N., Shapira, G., & Palino, T. (2017). *The Definitive Guide REAL-TIME DATA AND STREAM PROCESSING AT SCALE*. <https://doi.org/10.1017/CBO9781107415324.004>
- National Institute of Standards and Technology, N. (2013). Foundations for Innovation in Cyber-Physical Systems. *Workshop Report. National Institute of Standards and Technology*. <https://doi.org/10.1007/s13398-014-0173-7.2>
- Neisser, U. (2014). Cognitive Psychology. In *Cognitive Psychology*. Psychology Press. <https://doi.org/10.4324/9781315736174>
- Nissenbaum, H. (1996). Accountability in a computerized society. *Science and Engineering Ethics*, 2, 25–42.
- Nodarakis, N., Sioutas, S., Tsakalidis, A., & Tzimas, G. (2016). *Using Hadoop for Large Scale Analysis on Twitter: A Technical Report*. <http://arxiv.org/abs/1602.01248>
- OECD. (2020). *OECD Principles on Artificial Intelligence*. Organisation for Economic Co-Operation and Development. <http://www.oecd.org/going-digital/ai/principles/>
- Oleson, K. E., Billings, D. R., Kocsis, V., Chen, J. Y. C., & Hancock, P. A. (2011). Antecedents of trust in human-robot collaborations. *2011 IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support, CogSIMA 2011*, 175–178. <https://doi.org/10.1109/COGSIMA.2011.5753439>
- Ouaknine, A. (2018). *Review of Deep Learning Algorithms for Image Semantic Segmentation*. Medium. https://medium.com/@arthur_ouaknine/review-of-deep-learning-algorithms-for-image-semantic-segmentation-509a600f7b57
- Pääkkönen, P., & Pakkala, D. (2015). Reference Architecture and Classification of Technologies, Products and Services for Big Data Systems. *Big Data Research*, 2(4), 166–186. <https://doi.org/10.1016/j.bdr.2015.01.001>
- Pennachin, C., & Goertzel, B. (2007). Contemporary Approaches to Artificial General Intelligence. *Cognitive Technologies*, 8, 1–30. https://doi.org/10.1007/978-3-540-68677-4_1
- Pfeiffer, S., & Suphan, A. (2015). *The Labouring Capacity Index: Living Labour- der Universität Hohenheim ing Capacity and Experience as Resources on the Road to Industry 4.0*. <https://www.sabine-pfeiffer.de/files/downloads/2015-Pfeiffer-Suphan-EN.pdf>
- Phelps, E. S. (1997). *Rewarding Work: How to Restore Participation and Self-Support to Free Enterprise*.
- Pieters, W., & Cleeff, A. Van. (2009). THE PRECAUTIONARY PRINCIPLE IN A WORLD OF DIGITAL DEPENDENCIES. *Computer*, 42(6), 50–56. <https://doi.org/10.1109/MC.2009.203>
- Plank, B., & Hovy, D. (2015). *Personality Traits on Twitter—or—How to Get 1,500 Personality Tests in a Week*. 92–98. <https://doi.org/10.18653/v1/W15-2913>

- Plattform Industrie 4.0. (2016). *Structure of the Administration Shell Continuation of the Development of the Reference Model for the Industrie 4.0 Component*. <http://sf-eu.net/wp-content/uploads/2016/08/bmwi-2016-plattform-industrie-4.0-structure-of-the-administration-shell-en.pdf>
- Plutchik, R. (1982). A psychoevolutionary theory of emotions. *Social Science Information*, 21(4), 529–553.
- Poole, D. L., Mackworth, A., & Goebel, R. G. (1998). Computational Intelligence and Knowledge. *Computational Intelligence: A Logical Approach, Ci*, 1–22. <https://www.cs.ubc.ca/~poole/ci.html>
- Ramanathan, K. (2018). Enhancing Regional Architecture for Innovation to Promote the Transformation to Industry 4.0. In *Industry 4.0: Empowering ASEAN for the Circular Economy* (pp. 361–402).
- Rao, A. S., & Verweij, G. (2017). What's the real value of AI for your business and how can you capitalise? *PwC*, 27.
- Rauch, E., Matt, D. T., Brown, C. A., Towner, W., Vickery, A., & Santiteerakul, S. (2018). Transfer of industry 4.0 to small and medium sized enterprises. *Advances in Transdisciplinary Engineering*, 7(September), 63–71. <https://doi.org/10.3233/978-1-61499-898-3-63>
- Ritter, F. E., Tehranchi, F., & Oury, J. D. (2019). ACT-R: A cognitive architecture for modeling cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 10(3), 1–19. <https://doi.org/10.1002/wcs.1488>
- Roberts, D., & Johnson, R. (1997). Evolving frameworks: A pattern language for developing object-oriented frameworks. *Pattern Languages of Program Design*, 3(May 2014), 15. <https://doi.org/10.1.1.46.8767>
- Robinson, J. A. (1965). A Machine-Oriented Logic Based on the Resolution Principle. *Journal of the ACM (JACM)*, 12(1), 23–41. <https://doi.org/10.1145/321250.321253>
- Russell, S. J., & Norvig, P. (2012). *Künstliche Intelligenz: ein moderner Ansatz*. https://vowi.fsinf.at/images/b/bc/TU_Wien-Einführung_in_die_Künstliche_Intelligenz_VU_%28Eiter%2C_Tompits%29_-_Künstliche_Intelligenz-_Ein_moderner_Ansatz_%283.%2C_aktualisierte_Auflage%29.pdf
- Russell, S., & Norvig, P. (2010). Artificial Intelligence - A Modern Approach. In *Pearson*. <https://doi.org/10.1017/S0269888900007724>
- Salvucci, D. D. (2006). Modeling driver behavior in a cognitive architecture. *Human Factors*, 48(2), 362–380. <https://doi.org/10.1518/001872006777724417>
- Samaranayake, P., Ramanathan, K., & Laosirihongthong, T. (2018). Implementing industry 4.0 - A technological readiness perspective. *IEEE International Conference on Industrial Engineering and Engineering Management, 2017-Decem*, 529–533. <https://doi.org/10.1109/IEEM.2017.8289947>
- Samuel, A. L. (1959). Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal*, 3(3), 210–229. <https://doi.org/10.1147/rd.33.0210>
- Sanderson, K., & Andrews, G. (2006). Common Mental Disorders in the Workforce: Recent Findings From Descriptive and Social Epidemiology. In *Can J Psychiatry* (Vol. 51, Issue 2).
- Schaefer, K. E., Chen, J. Y. C., Szalma, J. L., & Hancock, P. A. (2016). A Meta-Analysis of Factors Influencing the Development of Trust in Automation: Implications for Understanding Autonomy in Future Systems. *Human Factors*, 58(3), 377–400. <https://doi.org/10.1177/0018720816634228>

- Schindler, K. (2020). *Visualisierung sowie Auswertung medizinischer Sensordaten unter Zuhilfenahme unterschiedlicher Algorithmen des maschinellen Lernens*. <https://epb.bibl.th-koeln.de/frontdoor/index/index/docId/1625>
- Schleipen, M. (2016). *OPC UA und AutomationML in der Industrie 4.0-Begriffswelt*. Informatik Aktuell. <https://www.informatik-aktuell.de/betrieb/netzwerke/opc-ua-und-automationml-in-der-industrie-40-begriffswelt.html>
- Schröder, C. (2016). The Challenges of Industry 4.0 for Small and Medium-sized Enterprises. *The Friedrich-Ebert-Stiftung*, 28.
- Searle, J. (2011). Watson Doesn't Know It Won on "Jeopardy!" *Wall Street Journal Online*, 257(43), 15. <https://www.wsj.com/articles/SB10001424052748703407304576154313126987674>
- Seeman, M. (1967). On the Personal Consequences of Alienation in Work. *American Sociological Review*, 32(2), 273. <https://doi.org/10.2307/2091817>
- Semmer, N. K. (2006). Job stress interventions and the organization of work. *Scandinavian Journal of Work, Environment and Health*, 32(6), 515–527. <https://doi.org/10.5271/sjweh.1056>
- Shepard, J. M. (1972). Alienation as a Process: Work as a Case in Point. *Sociological Quarterly*, 13(2), 161–173. <https://doi.org/10.1111/j.1533-8525.1972.tb00800.x>
- Siau, K., & Wang, W. (2018). *Building Trust in Artificial Intelligence, Machine Learning, and Robotics*. www.cutter.com
- Siau, K., & Wang, W. (2020). Artificial Intelligence (AI) Ethics. *Journal of Database Management*, 31(2), 74–87. <https://doi.org/10.4018/jdm.2020040105>
- Slama, D., Puhlmann, F., Morrish, J., & Bhatnagar, R. M. (2015). Enterprise IoT : Strategies & Best Practices for Connected Products & Services. In *O'Reilly Media*. <http://shop.oreilly.com/product/0636920039433.do>
- Smart Grid Coordination Group. (2012). *Smart Grid Reference Architecture* (Issue November). <ftp://ftp.cen.eu/EN/EuropeanStandardization/HotTopics/SmartGrids/Security.pdf>
- Smith, C., McGuire, B., Huang, T., & Yang, G. (2006). History of Artificial Intelligence. *Encyclopedia of Information Science and Technology, Second Edition, December*, 1759–1762. <https://doi.org/10.4018/978-1-60566-026-4.ch276>
- Sobocki, P., Lekander, I., Borgström, F., Ström, O., & Runeson, B. (2007). The economic burden of depression in Sweden from 1997 to 2005. *European Psychiatry*, 22(3), 146–152. <https://doi.org/10.1016/j.eurpsy.2006.10.006>
- Somers, S., & West, R. L. (2012). *Steering Control in a Flight Simulator Using ACT-R. July 2013*, 233–238.
- Sommer, L. (2015). Industrial revolution - Industry 4.0: Are German manufacturing SMEs the first victims of this revolution? *Journal of Industrial Engineering and Management*, 8(5), 1512–1532. <https://doi.org/10.3926/jiem.1470>
- Spath, D., Ganschar, O., Gerlach, S., Hämmerle, M., Krause, T., & Schlund, S. (2013). *Industrie 4.0 in Produktion, Automatisierung und Logistik*. Fraunhofer Verlag. <https://www.iao.fraunhofer.de/images/iao-news/produktionsarbeit-der-zukunft.pdf>
- Spectral Engines. (2018). *Industry 4.0 and how smart sensors make the difference*. Spectralengines.Com. <https://www.spectralengines.com/articles/industry-4-0-and-how-smart-sensors-make-the-difference>

- Speleta, J. (2019). *How Kubernetes Works*. DZone Cloud. https://dzone.com/articles/how-kubernetes-works?utm_source=dzone&utm_medium=article&utm_campaign=k8s-cluster
- Stewart, T. C., & West, R. L. (2006). Deconstructing ACT-R. *Proceedings of the Seventh International Conference on Cognitive Modeling*, 1(2), 298–303. <http://act-r.psy.cmu.edu/papers/641/stewartPaper.pdf>
- Stigler, S. M. (1989). Francis Galton's account of the invention of correlation. *Statistical Science*, 4(2), 73–79. <https://doi.org/10.1214/ss/1177012580>
- Stork, J., Eiben, A. E., & Bartz-Beielstein, T. (2018). A new taxonomy of continuous global optimization algorithms. In *arXiv* (pp. 1–24). arXiv. <https://doi.org/10.1007/s11047-020-09820-4>
- Strohschein, J., Fischbach, A., Bunte, A., Faeskorn-Woyke, H., Moriz, N., & Bartz-Beielstein, T. (2020). Cognitive Capabilities for the CAAI in Cyber-Physical Production Systems. *ArXiv E-Prints*. <http://arxiv.org/abs/2012.01823>
- Strohschein, J., Lara-Palma, A. M., & Faeskorn-Woyke, H. (2020). Employee technology acceptance of Industry 4.0 in SMEs. In M. Arias-Oliva, J. Pelegrín-Borondo, K. Murata, & A. M. Lara Palma (Eds.), *Societal Challenges in the Smart Society* (pp. 475–486). Universidad de La Rioja. <https://dialnet.unirioja.es/descarga/libro/769585.pdf>
- Strohschein, J., Lara Palma, A. M., & Faeskorn-Woyke, H. (2019). Detecting emotions in social media. a technological challenge to enhance youngest behavior. *28th AEDEM International Conference - Management in a Smart Society: Business and Technological Challenges*.
- Syafrudin, M., Fitriyani, N. L., Li, D., Alfian, G., Rhee, J., & Kang, Y. S. (2017). An open source-based real-time data processing architecture framework for manufacturing sustainability. *Sustainability (Switzerland)*, 9(11). <https://doi.org/10.3390/su9112139>
- Tasoulis, S. K., Vrahatis, A. G., Georgakopoulos, S. V., & Plagianakos, V. P. (2018). Real Time Sentiment Change Detection of Twitter Data Streams. *2018 IEEE (SMC) International Conference on Innovations in Intelligent Systems and Applications, INISTA 2018, Us*. <https://doi.org/10.1109/INISTA.2018.8466326>
- Tecuci, G. (2012). Artificial intelligence. *Wiley Interdisciplinary Reviews: Computational Statistics*, 4(2), 168–180. <https://doi.org/10.1002/wics.200>
- The Economist. (2017). The human cumulus: Artificial intelligence will create new kinds of work. *Economist*, 413(6051). <https://www.economist.com/business/2017/08/26/artificial-intelligence-will-create-new-kinds-of-work>
- The State Council of the People's Republic of China. (2018). *China to publish guideline on AI development*. http://english.www.gov.cn/state_council/ministries/2018/03/11/content_281476074108044.htm
- The White House. (2019). *Maintaining American Leadership in Artificial Intelligence*. <https://www.federalregister.gov/documents/2019/02/14/2019-02544/maintaining-american-leadership-in-artificial-intelligence>
- The White House. (2020). *American Artificial Intelligence Initiative: Year One Annual Report*. <https://www.whitehouse.gov/wp-content/uploads/2020/02/American-AI-Initiative-One-Year-Annual-Report.pdf>
- Tole, A. A. (2013). Big Data Challenges. In *Database Systems Journal: Vol. IV* (Issue 3).

- Turing, A. M. (1950). Computing machinery and intelligence. In *Machine Intelligence: Perspectives on the Computational Model* (Vol. 49, pp. 433–460). <https://doi.org/10.1093/mind/lix.236.433>
- Turnbull, J. (2014). *The Docker Book*.
- Urmersbach, B. (2019). *Anteile der Wirtschaftssektoren am BIP in Industrie- und Schwellenländern 2018* / Statista. Statista. <https://de.statista.com/statistik/daten/studie/37088/umfrage/anteile-der-wirtschaftssektoren-am-bip-ausgewaehlter-laender/>
- Van Brussel, H., Wyns, J., Valckenaers, P., Bongaerts, L., & Peeters, P. (1998). Reference architecture for holonic manufacturing systems: PROSA. *Computers in Industry*, 37(3), 255–274. [https://doi.org/10.1016/S0166-3615\(98\)00102-X](https://doi.org/10.1016/S0166-3615(98)00102-X)
- VDI. (2015). *Status Report: Reference Architecture Model Industrie 4.0 (RAMI4.0)*. https://www.zvei.org/fileadmin/user_upload/Presse_und_Medien/Publikationen/2016/januar/GMA_Status_Report__Reference_Architecture_Model_Industrie_4.0__RAMI_4.0_/GMA-Status-Report-RAMI-40-July-2015.pdf
- VDMA. (2016). Industrie 4.0 in practice: Solutions for industrial applications. *VDMA Industrie 4.0 Newsletter*, 1–58. https://industrie40.vdma.org/documents/4214230/26342484/Industrie_40_in_practice_2016_1529498623105.pdf/f5883098-521d-0b50-2a97-b1471ff13ace
- Verma, A., Pedrosa, L., Korupolu, M., Oppenheimer, D., Tune, E., & Wilkes, J. (2015). Large-scale cluster management at Google with Borg. *Proceedings of the 10th European Conference on Computer Systems, EuroSys 2015*. <https://doi.org/10.1145/2741948.2741964>
- Vidosav, D. (2014). Manufacturing Innovation and Horizon 2020—Developing and Implement „New Manufacturing“. *Proceedings in Manufacturing Systems*, 9(1), 3–8. http://icmas.eu/Journal_archive_files/Vol_9-Issue1_2014_PDF/3-8_MAJSTOROVIC.pdf
- Vincent, J. (2018). *DeepMind’s AI can detect over 50 eye diseases as accurately as a doctor*. The Verge. <https://www.theverge.com/2018/8/13/17670156/deepmind-ai-eye-disease-doctor-moorfields>
- Vogel, O., Arnold, I., Chughtai, A., & Kehrer, T. (2011). *Software Architecture*. Springer. <https://doi.org/10.1007/978-3-642-19736-9>
- Wang, W., & Siau, K. (2019). *Industry 4.0 Ethical and Moral Predicaments*. August.
- Watts, K. (2015). *Microservices Architecture: Deep exploration of Microservices*. CreateSpace Independent Publishing Platform. <https://dl.acm.org/doi/book/10.5555/2987980>
- Weise, T., Chiong, R., Lassig, J., Tang, K., Tsutsui, S., Chen, W., Michalewicz, Z., & Yao, X. (2014). Benchmarking optimization algorithms: An open source framework for the traveling salesman problem. *IEEE Computational Intelligence Magazine*, 9(3), 40–52. <https://doi.org/10.1109/MCI.2014.2326101>
- Whitener, E. M., Brodt, S. E., Korsgaard, M. A., & Werner, J. M. (1998). Managers as initiators of trust: An exchange relationship framework for understanding managerial trustworthy behavior. *Academy of Management Review*, 23(3), 513–530. <https://doi.org/10.5465/AMR.1998.926624>
- Winck, B. (2020). *Alphabet’s soaring stock just pushed it above a \$1 trillion market cap. Here are the 11 highest-valued public companies*. Business Insider. <https://www.businessinsider.de/international/highest-valued-public-companies-apple-aramco-biggest-market-cap-2020-1/?r=US&IR=T>
- Wischmann, S., Wangler, L., & Botthoff, A. (2015). *Industrie 4.0: Volks- und betriebswirtschaftliche*

- Faktoren für den Standort Deutschland.* <https://vdivde-it.de/system/files/pdfs/industrie-4.0-volks-und-betriebswirtschaftliche-faktoren-fuer-den-standort-deutschland.pdf>
- World Wide Web Consortium. (1999). *HTTP/1.1: Method Definitions.* <https://www.w3.org/Protocols/rfc2616/rfc2616-sec9.html>
- Wuest, T., Schmid, P., Lego, B., & Bowen, E. (2018). Overview of smart manufacturing in West Virginia. *Bureau of Business & Economic Research, West Virginia University.*
- Xing, B., Xiao, Y., Qin, Q. H., & Cui, H. (2018). Quality assessment of resistance spot welding process based on dynamic resistance signal and random forest based. *International Journal of Advanced Manufacturing Technology, 94*(1–4), 327–339. <https://doi.org/10.1007/s00170-017-0889-6>
- Zaefferer, M., & Rehbach, F. (2020). Continuous Optimization Benchmarks by Simulation. *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), 12269 LNCS*, 273–286. https://doi.org/10.1007/978-3-030-58112-1_19
- Zaharia, M., Chowdhury, M., Franklin, M. J., Shenker, S., & Stoica, I. (2010). Spark: Cluster computing with working sets. *2nd USENIX Workshop on Hot Topics in Cloud Computing, HotCloud 2010.*
- Zawadzki, P., & Zywicki, K. (2016). Smart product design and production control for effective mass customization in the industry 4.0 concept. *Management and Production Engineering Review, 7*(3), 105–112. <https://doi.org/10.1515/mper-2016-0030>
- ZVEI. (2016). *Beispiele zur Verwaltungsschale - Basisteil.*

DECLARATION

This work has not been submitted in substance for any other degree or award at this or any other university or place of learning, nor is being submitted concurrently in candidature for any degree or other award.

This thesis is being submitted in partial fulfillment of the requirements for the degree of PhD.

This thesis is the result of my own independent work and investigation, except where otherwise stated. Other sources are acknowledged by explicit references. The views expressed are my own.

Langenfeld, January 2021

Jan Strohschein